AUTHORIZED DOCUMENTATION

Novell Archive and Version Services 2.1 Administration Guide

Novell®
NetWare®

6.5 SP8

November 2009

www.novell.com

Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc., has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the Novell Legal Patents Web page (http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide

This guide describes how to configure and manage a Novell® Archive and Version Services 2.1 for NetWare® server to archive multiple interval-based versions of files for convenient access and retrieval by users. It is divided into the following sections:

Audience

This guide is intended for network administrators.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html (http://www.novell.com/documentation/feedback.html) and enter your comments there.

Documentation Updates

For the most recent version of the *NW 6.5 SP8: Novell Archive and Version Services 2.1 for Administration Guide*, visit the NetWare 6.5 SP8 Documentation Web site (http://www.novell.com/documentation/nw65).

Additional Documentation

For documentation on accessing and restoring archived file versions, see the OES 2 SP2: Novell Archive and Version Services User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/data/front.html).

For information about Novell Storage Services™, see the *NW 6.5 SP8: Novell Storage Services File System Administration Guide* (http://www.novell.com/documentation/nw65/stor_nss_lx_nw/data/front.html).

For information about Novell iManager, see the Novell iManager 2.7 Administration Guide (http://www.novell.com/documentation/imanager27) .

For information about in the command files for ArkManager, see Archive Definitions (http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/bsin3bl.html) in the *NDK: Virtual File Services for NetWare* (http://developer.novell.com/wiki/index.php/Virtual_File_Services_for_NetWare).

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux* and UNIX*, should use forward slashes as required by your software.

Overview of Archive and Version Services

1

Novell® Archive and Version Services 2.1 for NetWare® provides a convenient and cost-effective way for individual users to instantly restore previous versions of modified, renamed, or deleted network files. It helps to minimize the user's unproductive time and frees you to focus on other critical IT issues. The user simply views a list of previous interval-based versions of source files, selects the file needed, then recovers it. The user can recover any of the available versions. If users work in a collaborative environment, a user can determine which version to choose based on who modified a document and when.

Archive and Version Services is available for Novell Open Enterprise Server to archive user network files that are stored on Novell Storage Services™ volumes on NetWare 6.5 and later servers.

This section discusses the following:

1.1 Individual File Losses Impact Business

Most enterprises implement some type of data backup and recovery to prevent major data losses. Backups occur periodically to prevent catastrophic losses of data. Often, the files that individuals lose have a life cycle shorter than the major backup cycles. Until now, these data losses have been an unfortunate cost of doing business.

Recovery of a single file is not usually a simple process. Only the administrator can access the backup media to retrieve and recover the file. The user must know exactly when the file existed so that the administrator can find the right version of the file. Even after the file is recovered, the user must update the file with changes made between the time it was backed up until the time it was modified, deleted, or lost.

Individual losses of key data impact business. However, most enterprises leave prevention and recovery to the best practices and personal habits of users. In a typical network environment, users employ different techniques to ensure that they do not lose critical files. For example, some users manually save multiple versions of a file under different names. Others save the same version of a file in different locations. Some do both.

Despite precautions, almost every user has accidentally modified, lost, or deleted a key file. When problems occur, the user is left with two choices:

- Wait for the administrator to recover the file from backup media, if the file was backed up at all
- Painstakingly rebuild the file from a backup version or from scratch

Either solution negatively impacts business:

- It's inefficient. The user cannot access backup files without administrator action.
- It's inconvenient. The user must waste time re-creating materials.
- It can affect the enterprise's ability to meet business commitments. Time lost can impact the user's ability to meet milestones, thereby impacting delivery to other processes down the line.

1.2 Benefits of Archive and Version Services

Novell Archive and Version Services for NetWare provides benefits for the enterprise, IT administrators, and users.

Benefits for the Enterprise

Novell Archive and Version Services for NetWare offers two key benefits for the enterprise:

- It provides a lower cost of management for IT departments by allowing users to self-restore files from an archive of interval-based file versions.
- It provides a means to allow users to be more productive by allowing them to correct their own accidental deletions or file-modification mistakes.

Benefits for IT Administrators

For IT administrators, Novell Archive and Version Services for NetWare solves the problem of individual file recovery. No longer does the Help Desk need to deal with users asking for a particular file to be restored. This frees IT organizations to focus on more important solutions for the users and the company as a whole.

Benefits for Users

With Novell Archive and Version Services, the user controls the file recovery; there is no need to completely rebuild a file or to involve the IT department. Users can retrieve file versions from anywhere, at any time, using a Web browser and an active network or Internet connection.

Novell Archive and Version Services for NetWare offers many benefits for users:

- The versioning process is transparent to users until they need to retrieve a previous version of a file. Versioning does not affect how applications behave and requires no action on the part of the user.
- All security features and permissions of the source file are in effect for its file versions if the file versions are restored to an NSS volume on a NetWare 6.5 or later server. Users can also download a file version as a new file, without its prior rights and metadata, to other types of storage media, such as to their local workstations.
- Versioning supports collaborative work environments where groups of users can create and modify shared files. The archive server allows a work group to properly select previous versions of files they are working on, based on who modified a file and the time stamp of the version.
- Versioning works with any type of file from any type of application.

- Novell Archive and Version Services for NetWare uses Novell NetStorage to provide a Web-based interface to users for file version retrieval and restoration. Users can retrieve file versions from any workstation operating platform, including NetWare, Microsoft* Windows*, Apple* Macintosh*, Linux, and UNIX. All a user needs is a Web browser and an active network or Internet connection.
- Novell Archive and Version Services for NetWare provides the NSS File Version Utility for Windows clients. Users can retrieve file versions and restore them from compatible Windows workstations.

1.3 Key Concepts of Archive and Version Services

Novell Archive and Version Services for NetWare saves versions of user files at scheduled intervals, stores file versions in an archive database, and makes file versions available on demand to users. Novell Open Enterprise Server installs Novell Archive and Version Services by default as part of the NetWare Basic install. However, the service does not run until you configure and start it. For information, see Chapter 5, “Installing and Configuring New Archive Servers,” on page 47.

It is important to understand several key concepts and tools:

1.3.1 The Archive Server

The archive server runs Novell Archive and Version Services for NetWare, which includes the following services:

- Controlling the versioning process
- Providing the storage resources for the archive database
- Organizing and hosting the archive database
- Allowing users to search and restore file versions

An archive server and the volumes it serves reside in the same Novell eDirectory™ tree.

Novell Archive and Version Services supports several storage topologies, as shown in the following figure. Because file versions are transferred in decrypted format, the archive server should reside behind the corporate firewall. Any transfer of files during file versioning or restoration should occur over a secure connection such as a virtual private network (VPN). The volumes with files to be versioned can reside on local or remote servers and in single or clustered configurations.

For information, see "Planning for Archive and Version Services" on page 25.

Figure 1-1 *Example of Storage Topologies Supported by Novell Archive and Version Services*

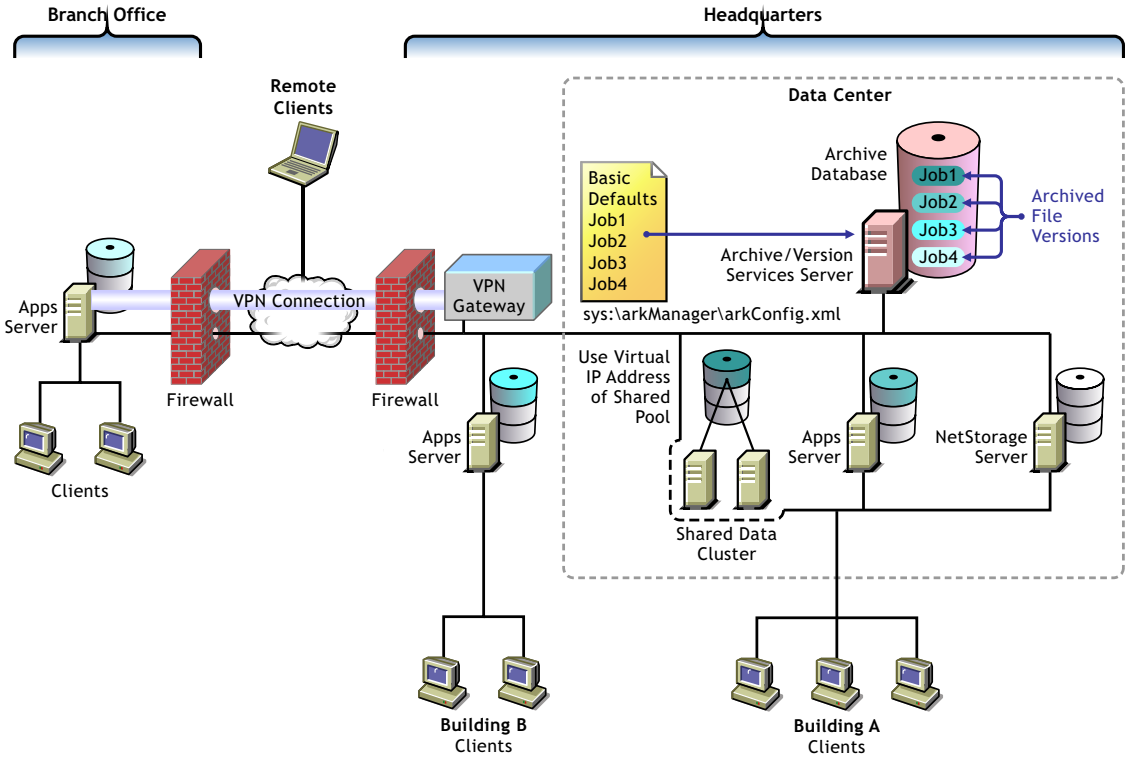

1.3.2 The Archive Database and ArkSQL

The archive database resides on the archive server. Novell Archive and Version Services 2.1 for NetWare uses a MySQL* server to organize and host file versions in the archive database. MySQL is an open source, structured query language (SQL) database. The `sys:\arkManager\ArkSQL.cnf` file is the configuration information for ArkManager's MySQL instance on the MySQL server. This guide refers to that instance of MySQL as ArkSQL.

1.3.3 Upgrade

Archive and Version Services 2.0 is upgraded to Archive and Version Services 2.1on upgrading the NetWare 6.5 SP 6 server to NetWare 6.5 SP7 server.

1.3.4 ArkManager and the ArkConfig.xml File

ArkManager is the software component in Archive and Version Services that manages the versioning, archiving, and file version management processes. For each archive server, you must configure the server's basic properties, optionally configure default job settings, and configure

properties for one or more individual jobs. Basic properties include details about the archive server and database. Default job properties specify the property settings to use instead of property settings specific to a particular job. An individual job specifies the property settings to use when archiving file versions for a specified volume that resides in the same eDirectory tree as the archive server. The ArkManager configuration information resides in the `sys:\arkManager\arkConfig.xml` file.

The Archive Versioning plug-in for Novell Archive and Version Services allows you to configure server, defaults, and job settings and to manage jobs. For information, see "Configuring Jobs in iManager" on page 59.

To configure versioning jobs using the `sys:\arkManager\arkConfig.xml` file, see "Configuring Jobs in ArkConfig" on page 77. For detailed information about ArkManager XML elements, see "XML Elements and Attributes for ArkConfig" on page 121.

1.3.5 Versioning Jobs

A versioning job captures copies of eligible files on a specified source volume at specified intervals. Eligible files are those that exist in the source volume at the time the volume is versioned, meet the general versioning criteria, and pass any administrator-specified filtering criteria. You can define only one job for a given source volume.

Each job identifies the settings for the following properties. For details, see Section 3.4, "Understanding Job Properties," on page 29.

Table 1-1 *Overview of Job Properties*

Property	Description
Name	The unique, administrator-specified job name that represents the relationship between the archive server and a given source volume. The job name persists for the life of the archive server and can represent only the specified source volume.
Source Server	The NetWare 6.5 or later server where the data to be versioned is located.
Source Volume	The NetWare 6.5 or later NSS volume where the data to be versioned is located. Each volume can be the target of only one versioning job.
Server Context	The Novell eDirectory context where the source server is located.
Snapshot Pool	The NetWare 6.5 or later NSS pool where the snapshots of the source volume are temporarily stored while file versions are written to the archive database.
Scheduled Interval	For a job, the elapsed time between the beginning of versioning processes.
Start Time and Schedule	The scheduled start time when the job's version process begins on one or more scheduled days of the week.
Delete Policy	Determines the lifetime of file versions in the archive.
Filter	Sets criteria to determine which files in the source volume are eligible for versioning.

1.3.6 Job Schedules

You must establish a schedule for each versioning job that meets users' requirements for file versioning, given limited storage and bandwidth resources. Versioning occurs for eligible files at scheduled intervals, called epochs. In Novell iManager, you can also manually pause versioning jobs and run jobs on demand, as needed.

A file's lifetime must span the end of an epoch to be versioned. Only files that exist when the versioning occurs are eligible to be versioned. If a user creates and deletes a file within the epoch, it cannot be versioned.

For more information, see Section 3.4.3, "Run Schedule," on page 31.

1.3.7 File Versions

File versions are actual copies of files taken at scheduled intervals, as determined by the administrator. No matter how many changes users makes to files during an epoch, only those eligible files that exist at the end of the epoch are saved.

Novell Archive and Version Services 2.1 for NetWare can use NSS pool snapshot technology to capture point-in-time copies of all files, even if the file is in use when the versioning process begins. If the snapshot option is not used, the versioning process captures only eligible files that are not deleted and not exclusively opened at the time.

User needs and limited storage and bandwidth resources are key considerations for setting the criteria to determine which files are eligible for versioning. Files can be filtered to include or exclude source files, according to their path, file extension, or filename patterns. If a user's files meet the filtering criteria, they are eligible for versioning. For information, see Section 3.4.5, "Filter," on page 32.

Users do not have direct control over which of their files get versioned, when the versioning occurs, or the state of their files when the epoch ends and the copy is made. Users can access files natively with the NSS File Version Utility on a Windows 2000/XP/2003 desktop, or they can access their file versions at any time and from anywhere using the NetStorage Archive function. For information, see the *OES2: Novell Archive and Version Services 2.1 User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front)*.

1.3.8 Delete Policy

The Delete Policy determines when and which of a job's file versions to automatically delete from the archive database. Versioned files can have a limited lifetime in the archive. You optionally configure a job's Delete Policy to set the maximum keep time and the maximum number of versions to retain. The Delete Policy can allow indefinite retention of at least one most recent versioned file.

For more information, see Section 3.4.4, "Delete Policy," on page 31.

1.3.9 Archive Versioning Plug-In for Novell iManager

After you configure your archive server and its versioning jobs, you can use the Archive Versioning plug-in for Novell iManager to manage those jobs. You can start and stop jobs, view a list of jobs, and view information about jobs, such as their current status, configuration details, and run schedules. You can also view the ArkManager log, which lists all normal, warning, and error messages for each job on the archive server.

For information, see Section 6.2, "Accessing the Archive Versioning Plug-In in iManager," on page 60.

1.3.10 NSS File Version Utility

The NSS File Version Utility provides convenient and direct access in a native Windows environment to archived versions of user files. The utility integrates with a Windows Explorer desktop to provide a Versions option, which allows users to view recent versions of their files in the archive database and restore the desired file. Users select the desired version of the file, then click *Restore* to download the file locally or to restore the file version to a network storage location.

The NSS File Version Utility is available on the welcome page on the *Novell Open Enterprise Server Products CD*. For information about using the utility, see the *OES2: Novell Archive and Version Services 2.1 User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front)*.

1.3.11 NetStorage Archive Function

File versions reside in the archive database on the archive server. Users can restore file versions from the archive database at any time from anywhere using the Archive function in Novell NetStorage. Using the NetStorage interface in the enterprise portal, a user views a list of available versions of a file. The user simply selects the previous version of the file, then clicks *Restore* to download the file version to a specified location where the user has the necessary permissions.

If a user restores the file version to a NetWare NSS storage location, the archive server recovers the file version and all the rights and metadata about the file. If a user opts to download the file version elsewhere, the file is saved as a new file, without the prior rights and associations.

For information, see the *OES2: Novell Archive and Version Services 2.1 User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front)*.

1.4 Scenarios for Using Archive and Version Services

Novell Archive and Version Services 2.1 for NetWare is a valuable asset in today's diversified workplace. This section discusses the following usage scenarios:

- Section 1.4.1, "Cross-Platform Work Environments," on page 20
- Section 1.4.2, "Group Collaboration," on page 20
- Section 1.4.3, "File and Directory Name Changes," on page 20
- Section 1.4.4, "Selective File Versioning," on page 20

1.4.1 Cross-Platform Work Environments

Novell Archive and Version Services for NetWare provides a Web-based interface to the archive database in Novell NetStorage with the Archive function. Users can retrieve file versions from any workstation operating platform, including NetWare, Microsoft Windows, Apple Macintosh, Linux, and UNIX. All a user needs is a Web browser and an active network or Internet connection. Users can restore the file version to a NetWare NSS storage location with all security and metadata intact, or download the file version as a new file to any other location where they have the necessary permissions to do so.

1.4.2 Group Collaboration

In a shared work group, a team works collectively to share information, create information, and process information. Files are regularly shuffled back and forth between users, and they are usually worked on by more than one person. Sections of a presentation are created by different people and either merged or are simply edited into an existing file that is passed around.

Research has shown that people tend to solve the problems associated with lack of versioning in PC network systems by attempting to do an ad hoc versioning system. Unfortunately, everyone does it a bit differently. Some put a version or date in the name of the file, such as `MarketAnalysisv3.ppt`. Others put dates in the filename, such as `MarketAnalysis2002Oct03.ppt`. Still others use file folders with versions or dates in the names of the folders. Most are not consistent with their techniques and many do not even try. It is especially troublesome when the files are shared, because not every personal scheme is alike. Even with these various methods, mistakes happen.

Novell Archive and Version Services for NetWare supports collaborative work environments. The archive interface allows a user to view previous versions and see instantly who was the modifier of each of the versions without opening file versions to attempt to ascertain who modified it.

For example, Tom and Alice worked together to prepare a presentation. Two weeks ago, Alice deleted some edits that Tom made to the file. Now, the team needs those edits back. Alice cannot recall when she deleted the edits. By going to the Web-based archive access, Alice can view previous file versions. The modifier of the file is listed next to each file version. Alice easily identifies the file version from about two weeks ago that shows up with Tom as the modifier of the file. She can view the file version or restore the file, as needed, to recover the lost modifications.

1.4.3 File and Directory Name Changes

File and directory names are likely to change during their lifetimes. Novell Archive and Version Services for NetWare supports file and directory renaming. It tracks changes to the filename, such as when a user renames a file at some point in the process of its creation and modification. It also tracks changes made to the file's subdirectory (or path) if it is changed.

1.4.4 Selective File Versioning

It is not desirable or practical to version every file in a volume. Novell Archive and Version Services for NetWare allows administrators to determine which files on their NetWare 6.5 or later servers get versioned and the versioning intervals on a per-volume basis.

For example, consider a source volume that comprises multiple directories: Users, Shared, and several directories for applications. Although the files in the Users and Shared directories change frequently, the files in the applications directories are fairly stable. Novell Archive and Version Services allows the administrator to exclude files in applications directories from versioning.

The administrator can selectively control the versioning frequency for each volume. For example, consider a Users volume with files that change intermittently throughout the day and a Shared volume with files that change at the end of each scheduled work shift. Novell Archive and Version Services allows the administrator to schedule 30-minute epochs for the Users volume and schedule the start time and subsequent epochs to coincide with shift changes for the Shared volume.

The administrator can selectively control the types of files versioned. For example, consider a Productivity volume that contains both Web development applications and user files. With Novell Archive and Version services, the administrator can specify that only files with extensions of `.doc`, `.html`, `.xml`, and `.pdf` be versioned in a Productivity volume.

1.5 What's Next

Use the following table to determine where to find information:

Table 1-2 *Possible Tasks*

To Perform This Task	Refer To
Planning your Novell Archive and Version Services 2.1 for NetWare implementation	Planning for Archive and Version Services (page 25)
Assessing your implementation plan against the prerequisites and guidelines	Prerequisites and Guidelines (page 37)
Setting up a new archive server	Installing and Configuring New Archive Servers (page 47)
Configuring the versioning jobs	Configuring Jobs in iManager (page 59)
Managing the archive server	Managing the Archive Server (page 99)
Managing the versioning jobs	Managing Jobs (page 85)
Managing security for archive services	Security Considerations for Archive and Version Services (page 119)

What's New

2

This section describes enhancements and additions to the Novell® Archive and Version Services 2.1 for NetWare 6.5 SP7 or later.

2.1 NetWare 6.5 SP 7

The features in this section have been added to Archive and Version Services 2.1 for OES 2 NetWare since the initial release of Netware 6.5 SP6.

2.1.1 Archive Versioning Plug-in to iManager 2.7

Archive Versioning plug-in has been redesigned to support the filter option. You can filter out files in the source volume that you do not want to version by using a series of Include and Exclude elements using iManager Plug-in. For more information on how to filter files, see Chapter 6, "Configuring Jobs in iManager," on page 59

Planning for Archive and Version Services

3

This section discusses how to plan and design a Novell® Archive and Version Services 2.1 for NetWare® solution to meet your business needs.

3.1 Assessing Your Versioning Needs

Before implementing Novell Archive and Version Services 2.1 for NetWare in your network environment, collect information about your system to assess your versioning needs. Ask and answer the following questions:

- **Users:** Which users would benefit from having an archived database of multiple historical versions of their network files? For example, in a university environment, you might provide versioning support for faculty and staff, but not for campus lab environments.

  Examine business and operational activities to identify and prioritize users’ versioning needs. Use this information to plan a strategic implementation of versioning throughout your network.

- **Volumes and Directories:** What volumes do these users use? How are the user directories organized on each volume?

  If you are versioning only selected users’ data on a volume, you can exclude all data, and then include each user’s data by path.

  If you archive files from an encrypted volume, the destination path for Archive Manager should also be on an encrypted volume. If the destination path is a nonencrypted volume, the versioned data is stored in a nonencrypted state.

  For any source volume, you can define only one versioning job. If you attempt to define multiple jobs for a volume on the same or different server, Archive and Version Services does not run as designed and data integrity in the archive database is compromised.

  - The archive server can be the same or different server as the source server.
  - A single archive server can archive many volumes with only one job defined per volume.
  - A given volume cannot have multiple jobs defined for it, even if the jobs run on different archive servers.

- **File Extensions:** What types of data within these volumes are candidates for versioning? For example, productivity files such as documents, spreadsheets, presentations, graphics, and other file types typical in your industry.

Many types of data do not need versioning and can quickly consume valuable space in your archive database. Avoid versioning system files, log files, and databases. In addition, do not version temporary files that have no value to users, such as temporary Internet files and temporary application files. Identify the extensions of temporary files to be excluded from the job.

- **Versioning Frequency:** How often does the identified data change on each volume? You should set the frequency to best satisfy the needs of the users. Look for patterns in frequency, such as the following:
  - Does frequency vary by types of data?
  - Does frequency vary by individual users?
  - Does frequency vary by categories of user, such as power users, support groups, or manufacturing?
  - Does frequency vary by physical location? For example, headquarters, branch offices, geographical regions, or departments.
  - Does frequency vary by time? For example, by time of day, day of the week, week in the month, month in the year, fiscal quarter, season (winter, spring, summer, or fall), or special events.
- **Storage:** What are the minimum initial storage capacity and scalability needs for your archive database? Determine this information by assessing your current storage requirements. For example:
  - How much storage space is consumed by each data type? For example, estimate the number of files and their average size.
  - What is the expected growth of data? For example, the additional number of megabytes or gigabytes of storage needed by users by month, quarter, or year.
  - What is the rate of growth in storage capacity by data type? For example, evaluate the percentage increase over time of a particular file type, such as documents, spreadsheets, presentations, graphics, or other file types typical of your industry.
- **Archive Database:** The potential size of the archive database depends on the combination of the following:
  - How many versioning jobs are defined for the server?
  - How often does each job run?
  - How many files are versioned on the source volume?
  - How often are the files modified and how does that correspond to the scheduled interval for the versioning epoch?
  - What is the delete policy for the job?
- **Topology:** Where in the enterprise network are the volumes stored? For example, a data center, a branch office, or work islands throughout the campus.
- **Delete Policy:** What type of retention policy for file versions do you need for each volume? Establishing a delete policy helps you control the storage capacity required by the archive database. For example:
  - How long do you need to retain versions?
  - What is the maximum number of versions to keep, given the number of source files versioned and their sizes?

- Do you need to keep at least one file version after its source file has been deleted?
- Do you need a complete copy of all the data to begin the archive database, or do you want the collection of file versions to grow only as data changes?

3.2 Designing the Archive and Version Services Topology

Novell Archive and Version Services 2.1 for NetWare supports several storage topologies, as shown in the following figure. In a typical enterprise, the archive server resides in the data center to provide convenient access to services for all users. The source volumes with files to be versioned can reside on local or remote servers in single or clustered configurations.

The number of archive servers you need depends on the number and size of file versions being archived and the distribution of users who need to access the versions. For example, if you have numerous volumes with many files changing at a high frequency, you might implement multiple archive servers. You should also consider the feasibility of backup for a server based on the anticipated size of the archive database and data.

For a geographically distributed enterprise, an archive server might reside in each regional data center. Key factors to consider are the speed, capacity, security, and cost of communications links between regions. Because the archive server transfers files nonencrypted across network connections, communications environments and connections must be secure.

Figure 3-1 *Example of Storage Topologies Supported by Novell Archive and Version Services*

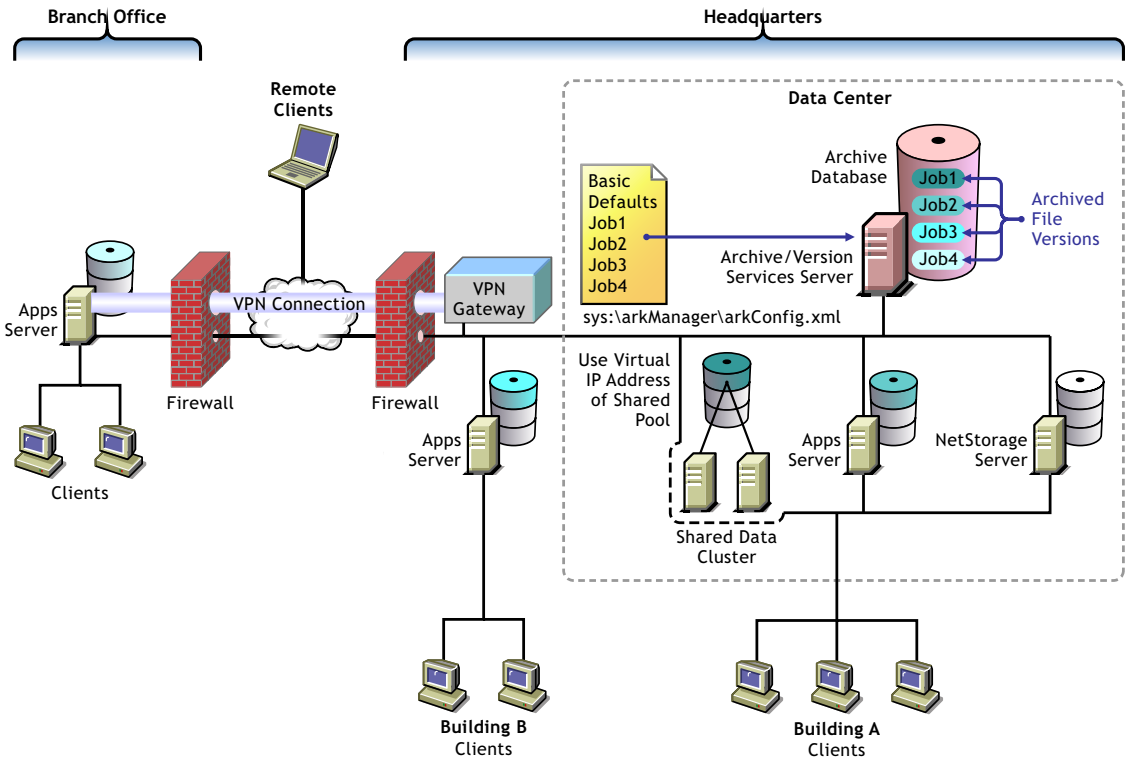

3.3 Understanding Archive Server Properties

The archive server properties define the basic information that applies to all jobs controlled by an archive server. They specify the authentication information for Novell eDirectory™ and the ArkSQL archive database, whether to display the log on the logger screen, and the storage location where the archive database (file versions) reside.

Configure the following server properties in the `sys:\arkManager\arkConfig.xml` file:

Table 3-1 *Description of Properties for the Archive Server*

Archive Server Properties	Description
Archive Path	Specifies the location of archive data in the NetWare 6.5 or later NSS archive volume. For example, `ark:\finance` or `archive:\users`.
Display Log	By default, the archive server records the error, warning, and normal messages for all of its jobs in the Archive Log (`sys:\arkManager\log`). If the Display Log option is enabled, the archive server prints the messages to the server's logger screen in addition to recording them in the log. You can view the logged messages in the Archive Versioning plug-in for iManager, the logger screen, or the archive log file.
eDirectory	Specifies authentication information about the Archive and Version Services administrator user. ◆ **User Name:** The eDirectory Common Name of the administrator user who has the appropriate rights to the original data location and to the archive data location. For example, `admin.servercontext` The archive server administrator user must have the file system Supervisor right to the archive server and to all servers being accessed by the selected archive server. The user must be in the same eDirectory tree as the archive server and the source servers. ◆ **Password:** The password of the archive server administrator for the specified eDirectory username. ◆ **Tree:** The eDirectory tree that contains the administrator user, the archive server, and the source servers.

Archive Server Properties	Description
Database	Specifies authentication information about the MySQL database for the archive server. ◆ **User Name:** The administrator user of MySQL. For example, root or mysqladmin. ◆ **Password:** The password of the MySQL administrator for the specified MySQL username. ◆ **Port:** The port used by the ArkManager instance of MySQL on the archive server. By default, Port 3306 is used. If Port 3306 is used by other services, other ports can be used instead, such as 3307 or 3308.

3.4 Understanding Job Properties

Job properties are the set of parameters used for individual job settings and default job settings. Individual job settings apply only to a single job. Default job settings can apply, in whole or in part, to any individual job defined for an archive server. Parameters used only for individual jobs include Name, Volume, and Stopped. Within a job definition, you can specify values for a given job property, or if you do not specify the value, its default value can be applied. In some cases, you can choose to specify no value, such as when a function is disabled.

Changing the property's value in an individual job's settings causes a different outcome than changing property's value for the default job settings. Modifying values in an individual job affects only the behavior of the individual job. Modifying values for the default job settings affects every job that uses the default values. All changes take effect the next time ArkManager runs.

Use the following properties to define jobs for your archive server:

- Section 3.4.1, "Job Information," on page 29
- Section 3.4.2, "Source Server Information," on page 30
- Section 3.4.3, "Run Schedule," on page 31
- Section 3.4.4, "Delete Policy," on page 31
- Section 3.4.5, "Filter," on page 32

3.4.1 Job Information

The Job Information identifies control information for jobs on an archive server. Each job has a unique name that represents a unique relationship between an archive server and a source volume.

IMPORTANT: You can define only one job per source volume.

The following table describes the job information properties:

Table 3-2 *Description of Job Properties for Job Information*

Job Property	Description
Name	The administrator-specified unique job name. For example, svr1_users or svr2_finance. A job name can be up to 64 ASCII characters. The job's name persists for the life of the archive server and represents the relationship between the archive server and a specific source volume. You cannot reuse the name for any other archive server and volume relationship. Specify the Name property only for individual jobs, not as a default value.
Stopped	Specify this property to define the job but leave it in a Stopped state until you manually activate the job. If the *Stopped* property is not used, the job starts, according to the Run Schedule settings, the next time ArkManager runs. Specify the *Stopped* property only for individual jobs, not as a default value.
Full Copy or No Full Copy	If Full Copy is selected, the archive server copies all files specified by the job from the source volume to the archive volume on a one-time, special version occurrence. Use this option to save at least one version of every file eligible for versioning. If No Full Copy is selected, the job begins at the next scheduled start time. Use this option to save versions of files eligible for versioning only as the files change during epochs.

3.4.2 Source Server Information

The Source Server information identifies the source server, the NSS volume that contains the information to be versioned, and the destination pool where temporary snapshot pools are stored during the versioning process.

Table 3-3 *Description of Job Properties for Source Server Information*

Job Property	Description
Server	The host name of the OES NetWare or NetWare 6.5 source server where the data to be versioned is located. For example, svr2. This server is in the same eDirectory tree as the archive server.
Volume	The name of the source NetWare 6.5 or OES NetWare NSS volume where the data to be versioned is located. For example, users or data. Specify the Volume property only for individual jobs, not as a default value.

Job Property	Description
Snapshot Pool	NSS pool snapshot technology allows all eligible files to be versioned at the end of the epoch, even exclusively open files. Specify the pool name of the destination pool for the snapshots. The snapshots are maintained temporarily at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, create a pool named `arksnaps` on the source server. If no snapshot pool name is specified or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume. In this case, eligible files that are exclusively opened at the time cannot be versioned, and eligible files that are deleted before the copy occurs cannot be archived.

3.4.3 Run Schedule

Run Schedule specifies when to start the job upon activation, and the frequency for running the job. To set the frequency, you must specify one of three scheduling options: Use Defaults, Scheduled Interval, or Scheduled Start Time.

Table 3-4 *Description of Job Properties for the Run Schedule*

Job Property	Description
Scheduled Interval	If it is used, the Scheduled Interval specifies the elapsed time between the beginning of versioning processes for a job. The units can be seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours. If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe the job skipping some versioning intervals, you can increase the interval between versions or reduce the amount of data to be versioned by setting Filter properties.
Scheduled Start Time	If it is used, the Scheduled Start Time specifies the start time when the job's version process begins and one or more days of the week to run the job. ◆ **Start:** Specify the start time in hours and minutes. ◆ **Daily:** Specify the days of the week to run the job. For example, specify All days for a daily job run, or specify one or more days of the week you want the job to run. Choices include Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday, and All.

3.4.4 Delete Policy

The Delete Policy determines the retention of file versions by age or by number of versions.

Table 3-5 *Description of Job Properties for the Delete Policy*

Job Property	Description
Delete Interval	Specifies the schedule for deleting file versions, if they are eligible for deletion. The interval represents the amount of time to wait from the time a Delete process ends until another Delete process begins. If a value is not specified, 24 hours is the default interval. The time involved in deleting file versions varies with the amount and complexity of data stored in the archive server. The Delete Schedule operates separately from the Version Schedule. For example, suppose you set the Delete Schedule to 1 hour. When you activate the job, the Delete Schedule begins an interval timer. After 1 hour elapses, the Delete process runs. The timer is inactive while the process runs. When the delete process ends, the interval timer begins again. The process repeats until the job is deactivated.
Maximum Keep	Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the Maximum Keep Versions and Maximum Keep Time to zero values, the Delete Policy function does not run. ◆ **Time:** The maximum time that a file version is maintained in the archive. Specify the time in whole numbers for seconds, minutes, hours, or days. ◆ **Versions:** Specify the maximum number of versions of each file to keep in the archive. When the number of versions exceeds this integer value, the oldest version is deleted. ◆ **Keep Current Copy:** By default, at least one file version of current files remains in the database, even if the Maximum Keep Time elapses. If Keep Current Copy is enabled, the archive keeps an existing file version as long as its source file is current on the source volume, beyond the Maximum Keep Time. After the user deletes the current source file, the deletion is noted at the next scheduled epoch. If the file version's age is within the Maximum Keep Time, the archive database retains a copy of the file version until its Maximum Keep Time elapses. When the file version's age exceeds the Maximum Keep Time, the archive deletes the file version at the next scheduled delete time. If Keep Current Copy is disabled, the archive deletes the file version when the Maximum Keep Time elapses.

3.4.5 Filter

The Filter information determines what data in the source volume gets versioned. You can combine the filters for the individual job with filters for the default job settings.

You can filter files in the source volume that you do not want to version by using a series of Include and Exclude elements under the Defaults element or an individual Job element in the `sys:\arkManager\arkConfig.xml` file. The order of the Include and Exclude elements determines what data is eligible for versioning. Make sure the order is adequate to achieve the desired filtering outcome.

Filtering is optional, but you should avoid versioning volumes that contain system software. Exclude system files and file types that change constantly, such as log files and databases. You can also exclude nonessential file types such as MP3 and temporary files such as Internet files. Identify the file types your applications use as intermediate saves for open files, such as the TMP files for Microsoft Word, and set up filters to exclude that file extension from versioning.

Extremely large files, such as database files and ISO image files, take a long time to be copied into the archive, which can potentially block other requests to access the database. You cannot filter files by file size, but you can modify database settings or distribute data to lessen the impact of versioning large files. Another option is to separate larger files into one or more separate volumes, and then create a job for each source volume. Schedule the jobs to run in off-peak hours.

For example, if your files are potentially very big, such as several hundred megabytes to gigabytes in size, you might need to increase the time that queries wait to access a locked MySQL database before time-out. For example, set the `MySQL Lock Wait Timeout` variable (`innodb_lock_wait_timout`) in the `sys:\arkManager\arkSQL.cnf` file to more than the default 50 seconds.

Table 3-6 *Description of Properties for the Filter*

Job Property	Description
Path	Specifies the relative path of directories in the specified source volume that you want to include or exclude in the versioning process. For example: `<path>\log</path>`
Extension	Specifies the extension of files in the specified source volume that you want to include or exclude in the versioning process. Use the preceding dot followed by the characters of the file extension. For example: `<extension>.mp3</extension>`
Pattern	Specifies the regular expression pattern to match for files that you want to include or exclude in the versioning process. For example, to specify a pattern for files that start with the letter `a`: `<pattern>.*\\a.*</pattern>`
Wildcard	Specifies the wildcard pattern to match for files that you want to include or exclude in the versioning process. For example: `<wildcard>a*tmp</wildcard>`

Include and Exclude Elements

You can filter out files in the source volume that you do not want to version by using a series of Include and Exclude elements. The child elements within each Include or Exclude element can contain multiple Path, Extension, and Pattern elements, in whatever order is needed to determine what data is eligible for versioning.

By default, all data in the volume is included. The first step in filtering is to exclude everything. For example:

```
<exclude>
  <path>\</path>
</exclude>
```

Next, add back in the paths, file types, and patterns for files you want to version. The latter include or exclude elements override previous include or exclude elements.

Pattern Elements

The regular-expression parser used for the `<pattern>` tag does not support the following regular-expression constructs in PERL 5:

- The conditional constructs `(?{X})` and `(?(condition)X|Y)`
- The embedded code constructs `(?{code})` and `(??{code})`
- The embedded comment syntax `(?#comment)`
- The preprocessing operations `\l`, `\u`, `\L`, and `\U`

The regular-expression parser used for the `<pattern>` tag supports the following regular-expression constructs, which PERL 5 does not:

- Possessive quantifiers, which match as much as they can and do not back off, even when doing so would allow the overall match to succeed
- Character-class union and intersection

  Character classes can appear within other character classes, and can be composed by the union operator (implicit) and the intersection operator (`&&`). The union operator denotes a class that contains every character that is in at least one of its operand classes. The intersection operator denotes a class that contains every character that is in both of its operand classes.

  The precedence of character-class operators is as follows, from highest (1) to lowest (5):

  1) Literal escape `\x`
  2) Grouping `[...]`
  3) Range a-z
  4) Union `[a-e][i-u]`
  5) Intersection `[a-z&&[aeiou]]`

Other notable differences from PERL-based regular expressions are shown in the following table.

Table 3-7 *Comparison of Supported Regular Expressions and PERL-Based Regular Expressions*

Supported Regular Expressions	PERL-Based Regular Expressions
Octal escapes must always begin with a zero. `\1` through `\9` are always interpreted as back references, and a larger number is accepted as a back reference if at least that many subexpressions exist at that point in the regular expression; otherwise, the parser drops digits until the number is smaller or equal to the existing number of groups or it is one digit.	`\1` through `\9` are always interpreted as back references; a backslash-escaped number greater than 9 is treated as a back reference if at least that many subexpressions exist; otherwise, it is interpreted, if possible, as an octal escape.
It is implicit that repeated invocations of the find method resume where the last match left off, unless the matcher is reset.	PERL uses the g flag to request a match that resumes where the last match left off.
Embedded flags always take effect at the point where they appear, whether they are at the top level or within a group; in the latter case, flags are restored at the end of the group just as in PERL.	In PERL, embedded flags at the top level of an expression affect the whole expression.
The defined class accepts dangling brackets but is strict about dangling metacharacters like `+`, `?` and `*`, and throws a PatternSyntaxException if it encounters them.	PERL is forgiving about malformed matching constructs, as in the expression `*a`, as well as dangling brackets, as in the expression `abc]`, and treats them as literals.

For more information, consult a programming textbook or search the Internet for a reference that discusses the behavior of regular expression constructs.

Wildcard Elements

A wildcard functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to include or exclude files that start with d of type `.sxi`:

```
<wildcard>d*sxi</wildcard>
```

3.5 What's Next

Before you deploy Novell Archive and Version Services 2.1 for NetWare, make sure your deployment plan satisfies the following:

- Chapter 11, "Coexistence and Migration Issues for Archive and Version Services," on page 115
- Chapter 12, "Security Considerations for Archive and Version Services," on page 119
- Chapter 4, "Prerequisites and Guidelines," on page 37

Prerequisites and Guidelines

4

This section discusses prerequisites and guidelines for designing your Novell® Archive and Version Services 2.1 for NetWare® system:

4.1 Network Architecture Prerequisites and Guidelines

Make sure your network meets the following prerequisites:

- Because the archive server moves data in decrypted format, the archive server should be located behind the corporate firewall. If you are versioning data on remote servers, use a virtual private network (VPN) connection between the two.
- Determine how much bandwidth is required to support versioning traffic, given the following criteria:
  - How much data is transferred as file versions during each epoch
  - When the scheduled epochs occur
  - Where in the network the source volumes are located (the path between the archive server and the source server)
  - Peak traffic flow based on seasonal variations in productivity

4.2 Server Prerequisites and Guidelines

Make sure your archive server and source servers meet the following prerequisites and guidelines:

- The archive server can be the same or different server as the source server.

- One archive server can have only one job defined for any one volume.
- A single archive server can archive many volumes with only one job defined per volume.
- A given volume can have only one job defined for it.
- The archive server and the source server can be the same machine or different machines.
- Archive and Version Services is compatible with OES NetWare and NetWare 6.5. An archive server and its source servers can run either operating system in a mixed environment. Make sure the servers are running the latest upgrades.
- Make sure your server meets the system and software requirements for OES NetWare. For information, see the OES 2: NetWare Installation Guide (http://www.novell.com/documentation/oes2/inst_oes_nw/index.html?page=/documentation/oes2/inst_oes_nw/data/front.html#front).
- ArkManager uses UTF-8 encoding for filenames. Either the Language Code pages of the NCP™ client and server must be the same, or you must use a current version of the Novell Client™ to provide UTF-8 support and enable the UTF-8 service in the client. CIFS clients have included UTF-8 encoding support since NetWare 6.0 SP 2 and NetWare 6.5. For information, see Section 9.11, “Enabling UTF-8 Encoding Support for Clients,” on page 104.
- Determine how much storage space you need for the archive database, given the following criteria:
  - How much data needs to be stored immediately by the versioning job
  - How much data there is to version (the physical storage requirements by versioning job)
  - How frequently the data must be versioned (the length of time in an epoch by versioning job)
  - How long you need to store the versions (maximum keep time and maximum versions to keep by versioning job)

4.3 Novell eDirectory Prerequisites and Guidelines

Install and configure Novell eDirectory™ for your network, including the tree and server context where you want your archive server to reside. Make sure it is active and running properly. For information, see the Novell eDirectory 8.8 Administration Guide (http://www.novell.com/documentation/beta/edir88/edir88/index.html?page=/documentation/beta/edir88/edir88/data/a2iii88.html#a2iii88) .

Your archive server must be located in the same eDirectory tree as the servers that host its source volumes.

4.4 Novell iManager 2.7 Prerequisites

Install Novell iManager 2.7, making sure to install the Archive Versioning plug-in for OES SP1. For more information, see the Novell iManager 2.7 Administration Guide (http://www.novell.com/documentation/imanager27/index.html?page=/documentation/imanager27/imanager_admin_27/data/bsxrjzp.html#bsxrjzp).

4.5 Storage Media Prerequisites and Guidelines

Make sure your archive volume and source volumes meet the following prerequisites and guidelines:

- The source and archive data volumes must be Novell Storage Services™ (NSS) data volumes that use the OES NetWare and NetWare 6.5 media format.

  You can create the NSS volumes on an OES NetWare server, an OES Linux server, or a NetWare 6.5 server. If the NSS volume is on an OES Linux server, you must move it to a cross-compatible OES NetWare server before it can participate in your Archive and Version Services solution.

- If you are upgrading an existing archive server:
  - Back up your existing archive database and the archive data before you begin the upgrade to OES 2 NetWare. For information, see Section 9.8, "Backing Up the Archive Database," on page 103 and Section 9.9, "Backing Up the Archive Data," on page 103.
  - Verify that ArkManager is not running when you begin the upgrade. For information, see Section 9.2, "Stopping ArkManager," on page 99.
  - After the upgrade, continue to use the same NSS pool, volume, and directory for your upgraded archive database and archive data.
- If the volumes with files you want to version reside on NetWare 6.0 and earlier servers, perform the following operating system upgrades, as appropriate:
  - Upgrade the source server to OES 2 NetWare or NetWare 6.5 SP 7. For information, see the OES 2: NetWare Installation Guide (http://www.novell.com/documentation/oes2/inst_oes_nw/index.html?page=/documentation/oes2/inst_oes_nw/data/front.html#front) .
  - If you are upgrading from NetWare 6.0, your NSS media format upgrades in the background over a period of up to 21 days. The format update process is accelerated automatically if the volume begins to participate as a source volume before the conversion is complete; the upgrade requires no outside action on your part.
  - If you are upgrading from NetWare 5.*x*, you must upgrade your NetWare 5.*x* NSS volumes and Traditional volumes. For information, see Upgrading Legacy NSS Volumes and NetWare Traditional Volumes to NSS on NetWare Volumes in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).
- Create an NSS storage pool and volume to use exclusively for your archive database and archive data that meets the following guidelines:
  - Do not place the archive volume in the system (`sys`) pool.
  - For each archive server, use the information from your planning session to design the physical storage media for its archive database.
  - Make sure your implementation allocates enough space to meet your immediate archive needs and is scalable to accommodate future growth.
  - The NSS volume you use for your archive database cannot contain data that you plan to archive. The archive volume and the source volume cannot be the same volume.
  - If you plan to version files from an encrypted source volume, the archive volume should also be an encrypted volume so that data continues to be protected in the archive database. For information about Encrypted Volume Support, see Managing Encrypted NSS Volumes

in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

- The archive volume and the source volume should reside in separate storage pools, but they can be in the same pool.
- The pool for the archive volume and the pool for the source volumes should use different storage devices, but this is not required.
- You should store the archive database and archive data in a directory in the archive volume for easy backup. The archive data is always stored in subdirectories at the base of the archive path.

4.6 MySQL Prerequisites and Guidelines

Before you install the Novell Archive and Version Services 2.1 instance of MySQL on your MySQL server:

- If this is a new install, create the pool, volume, and directory where you want to store the archive database and archive data for Novell Archive and Version Services 2.1. For information, see Section 5.3.4, “Creating an Archive Volume,” on page 51.
- If this is an upgrade, continue to use the existing pool, volume, and directory for the archive database and archive data.
- Consider security issues for MySQL. If you are currently using a MySQL server, modify its settings to ensure a secure installation. If you are setting up a new MySQL server, make sure to choose the Secure Installation options.

  By default, the MySQL installation does the following:

  - Allows the MySQL database to be created without specifying a root password
  - Allows a root user to connect to the MySQL database from the local host or remotely
  - Creates an anonymous user and allows that user to connect locally or remotely
  - Allows the anonymous user to perform any function on any databases named “test” or that begin with “test_”
  - Creates an initial test database.

  With the default installation instead of a secure installation, any local user can connect without a password and be treated as the anonymous user. This creates a high security risk in your production environment. You are strongly encouraged to use a secure installation of MySQL for production environments.

For information about MySQL, see the *MySQL Reference Manual* (http://dev.mysql.com/doc/mysql/en/index.html).

4.7 ArkSQL Guidelines

Archive and Version Services uses MySQL with InnoDB, which manages all tables within one tablespace. The potential size of the tablespace in the archive database depends on the combination of the following:

- Number of versioning jobs defined
- The scheduled frequency of each job

- The expected number of source files and their modification frequency for each job
- The delete policy for file versions

By default, the initial amount of space reserved for the archive database tablespace is 400 MB. The size of the reserved space is automatically extended by 8 MB every time the tablespace is filled.

You can specify a larger or smaller initial tablespace by modifying the `innodb_data_file_path` parameter in the `sys:\arkManager\arkSQL.cnf` file. The minimum tablespace size is 10 MB, as governed by the InnoDB constraints. The reserved size must not exceed the volume size. The maximum tablespace size on NetWare is 8 TB because the maximum size of an NSS volume is 8 TB. However, because your archive database and archive data files share the same volume, the practical limit of the tablespace is much smaller than 8 TB.

For information about configuring ArkSQL, see Section 9.5, "Modifying the ArkSQL Settings," on page 101.

You can also extend the existing tablespace manually by adding new components (raw partition or regular file). For information, see "The InnoDB Storage Engine" (http://dev.mysql.com/doc/mysql/en/InnoDB.html) in the *MySQL Reference Manual* (http://dev.mysql.com/doc/mysql/en/index.html).

4.8 Fault Tolerance Guidelines

If your data is critical, you can design fault tolerant and high availability solutions for Novell Archive and Version Services, including multiple connection channels, software RAID devices, and cluster solutions. These solutions are optional.

4.8.1 Multiple Connection Channels for Storage Devices

Multiple connection channels can help ensure fault-tolerant connectivity between the archive server and the devices containing the archive volume.

For information, see Managing Multiple I/O to Devices (NetWare) in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front) .

4.8.2 Software RAID Devices

Software RAIDs can improve read/write performance and ensure data protection. NSS file systems support software RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping with parity). All software RAID devices can improve file access performance, but only RAID 1 and RAID 5 also provide data protection to your storage media solution. To add fault tolerance your archive volume, you can set up the device you plan to use for your NSS pool as a RAID-1 device, a RAID-5 device, or a RAID-10 device.

For information, see Managing Software RAID Devices in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

4.8.3 Server Clusters

A Novell Cluster Services™ solution improves service availability. In an active/passive cluster configuration, one server is active and any other server nodes act as hot-standby servers. If the active server goes down, the Cluster Services software handles the failover to the next available server in the cluster. If the server running Archive and Version Services goes down, Cluster Services ensures that the archive services and database remains available to versioning processes and to users who need to retrieve file versions.

Before you configure Novell Archive and Version Services in a cluster solution, you must install OES Cluster Services 1.7 for NetWare. OES includes Cluster Services and licenses for two cluster nodes. An active/passive, two-node NetWare cluster is the basic fault-tolerant solution.

If you plan to set up the archive server in a clustered configuration, you must create a shared NSS volume as your archive volume. If you are combining a clustered configuration and using software RAID devices, make sure the software RAID devices are sharable for clustering, then assign them as devices in the cluster-enabled NSS pool you create for your archive volume.

For information, see the OES 2: Novell Cluster Services 1.8.4 for NetWare Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs).

4.9 Job Guidelines

One archive server can have only one job defined for any one volume. However, the archive server can run jobs for multiple source volumes.

WARNING: Define only one job per volume. If you attempt to define multiple jobs for a volume, ArkManager does not run as designed and data integrity in the archive database is compromised.

Jobs should focus on versioning files from productivity applications. Filtering is optional, but you should avoid versioning volumes that contain system software. Exclude system files and file types that change constantly, such as log files and databases. You might also want to exclude nonessential file types such as MP3 and temporary files such as Internet files. Identify the file types your applications use as intermediate saves for open files, such as the TMP files for Microsoft Word, and set up filters to exclude that file extension from versioning.

Extremely large files, such as database files and ISO image files, take a long time to be copied into the archive, which can potentially block other requests to access the database. You cannot filter files by file size, but you can modify database settings or distribute data to lessen the impact of versioning large files. If your files are potentially very big, such as several hundred megabytes to gigabytes in size, you might need to increase the time that queries wait to access a locked MySQL database before time-out. For example, set the `MySQL Lock Wait Timeout` variable (`innodb_lock_wait_timout`) in the `sys:\arkManager\arkSQL.cnf` file to more than the default 50 seconds. Another option is to separate larger files into one or more separate volumes, and then create a job for each source volume. Schedule the jobs to run in off-peak hours.

A job name must be unique to the archive server. To see job names that are currently in use, see Section 8.1, “Viewing a Jobs Report,” on page 85. However, after you delete a job, it no longer appears in this list.

4.10 Schedule Guidelines

You should allow enough time for one version job to be captured before beginning another epoch for that job. If a versioning process cannot finish before the next process is scheduled to begin, the archive server continues to version the job and skips other scheduled processes until the job finishes.

For example, suppose that you configure a versioning job with a one-minute epoch for a source volume that contains an amount of data that takes 5.2 minutes to archive. The version process begins at time 0. At each 1-minute interval thereafter, the next scheduled versioning checks to see if the job is already running. If it is, the job does not begin. In this example, the scheduled versioning processes do not occur at 1, 2, 3, 4, and 5 minutes. The next versioning process begins at 6 minutes. Although the user expects to have file versions available at 1-minute intervals, the file versions are at 6-minute intervals.

To avoid such problems, consider the following configuration alternatives:

- Set epochs realistically, based on the amount of data to be versioned from the source volume and the bandwidth of the connection between the source server and the archive server.
- Filter the data on the source volume to exclude unnecessary data from being versioned.
- Divide data on the source volume to create multiple volumes. Then you can configure multiple jobs to run concurrently on the same or different archive server.

4.11 Prerequisites for Using the NSS File Version Utility

Users can search for file versions and restore them with the NSS File Version Utility on their Microsoft Windows 2000/XP/2003 clients, with either NCP or CIFS support. The utility integrates with the Windows desktop to allow users to natively restore previous versions of their current, renamed, or deleted files that are stored in the archive database.

- ❑ Set up NCP or CIFS support for the archive server.
- ❑ Distribute and install the NSS File Version Utility on compatible user workstations.

  You can find the NSS File Version Utility installation file (`nwver.exe`) in the *Novell Open Enterprise Server Products* CD in the `\tools\nsstools` directory.

For information about installing and using the NSS File Version Utility, see the OES 2: Novell Archive and Version Services 2.1 User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front).

4.12 Prerequisites for Using the NetStorage Archive Function

Users can search for file versions and restore them with Novell NetStorage from anywhere, at any time, using a compatible Web browser and a network or Internet connection.

- ❑ Install a Novell NetStorage server.
- ❑ Configure Archive access to NetStorage services.

For information about installing and configuring Novell NetStorage, see the OES 2: NetStorage for NetWare Administration Guide (http://www.novell.com/documentation/oes2/file_netstor_nw/index.html?page=/documentation/oes2/file_netstor_nw/data/h9izvdye.html#h9izvdye).

For information about using NetStorage and the Archive function, see the OES 2: Novell Archive and Version Services 2.1 User Guide (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front) .

4.13 Prerequisites and Guidelines for Retrieving File Versions

Previous versions of files can exist in the archive database only after the files meet the following prerequisites:

- You install and configure Novell Archive and Version Services 2.1 on your OES NetWare server.
- You configure at least one job to version OES NetWare or NetWare 6.5 NSS volumes where user data is stored. These NSS data volumes can reside on an OES NetWare server or NetWare 6.5 server.
- You set criteria to determine which files are eligible for versioning. Files can be included or excluded according to their path, file extension, or filename patterns. If files on a specified volume meet the resultant criteria, they are eligible for versioning.
- Versioning occurs for eligible files at scheduled intervals, called epochs. You can also use iManager to pause jobs and run jobs on demand, as needed.
- A file's lifetime must span the end of at least one epoch to be versioned. Only files that exist when the versioning occurs are versioned.

  If a user creates a file after an epoch begins, and deletes the file before the end of the epoch, the file cannot be versioned. If a user creates a file that spans the end of a scheduled epoch, but deletes the file before the scheduled job actually copies the file to the archive, the file is not versioned.
- It does not matter how much or how often a user changes an eligible file during an epoch. The versioning process captures the file in whatever state it is in at the end of the epoch. A user does not have direct control over which files are versioned, when the versioning occurs, or what the state of any file is when the epoch ends and the file versions are copied to the archive database.
- Versioned files typically have a limited lifetime in the archive. You configure a Delete Policy that sets the maximum keep time and the maximum number of versions to retain. The Delete Policy for volumes can be configured to allow indefinite retention of at least one most recent file version of a current file.
- To retrieve file versions using the NSS File Version Utility, users must download and install the NSS File Version utility on their Windows 2000/XP/2003 workstations.
- To retrieve file versions using Novell NetStorage, users must have a workstation with a compatible Web browser and a network or Internet connection to the Novell NetStorage server.

4.14 Guidelines for Availability of File Versions

After your system meets the Prerequisites and Guidelines for Retrieving File Versions, there are several reasons that previous versions of a current, renamed, or deleted file might not exist in the archive database:

File Versions Are No Longer Supported for a Volume

If you suspend or delete a versioning job for a volume, its scheduled jobs do not run and no new versions are saved to the archive. For suspended jobs, users can retrieve available file versions from the archive. For deleted jobs, the file versions remain in the database but users cannot access them until a job with the same name is added back into the `sys:\arkManager\arkConfig.xml` file. For suspended jobs, the job's Delete Policy continues to be applied to existing versions in the archive. Eventually, the file versions meet the criteria for deletion defined in the job's Delete Policy.

If Keep Current Copy is enabled for the job's Delete Policy, at least one file version of a file remains in the database if its source file is current on the source volume, even if the Maximum Keep Time elapses.

If Keep Current Copy is disabled, the archive deletes the file version when the Maximum Keep Time elapses.

A File's Lifetime Does Not Span the End of an Epoch

A file is versioned only if it exists at the end of an epoch. If scheduled versioning processes are paused or delayed for a period of time that exceed the lifetime of a given file, its file version might not be captured, even if the file exists for a length of time that exceeds the epoch. Although the frequency of the epoch usually is less than the typical lifetime of your files, there might be occasions when the versioning time is extended. For example, if the time it takes to save versions of all the files eligible for versioning overlaps one or more scheduled epochs, some scheduled epochs might be skipped until the current process ends and the next scheduled start time occurs.

To avoid this problem:

- Schedule a job's epoch for a period of time shorter than a typical file's lifetime for the volume, while allowing enough time to capture file versions.
- Enable the Pool Snapshot option for capturing file versions.
- Minimize the use and length of job pauses.

A File Is Open When Versioning Occurs

Novell Archive and Version Services 2.1 leverages NSS pool snapshot technology to save point-in-time versions of all files, including open ones. If you disable the Pool Snapshot option, the files that are eligible for versioning are copied directly from the source volume. Exclusively open files cannot be versioned and data might be inconsistent.

To avoid this problem, make sure to enable the Pool Snapshot option for the archive server.

A File Is Deleted Before A Version Can Be Saved

If the Pool Snapshot option is disabled for your archive server, file versions are saved from the volume to the archive, which takes time. If a file is created during an epoch and exists at the end of an epoch, but a user deletes it before a version can be copied to the archive database, that file version

is not saved to the archive database. If the Pool Snapshot option is enabled, versions of all files, even exclusively open ones, are captured by the snapshot and a point-in-time version of each file can be saved to the archive.

To avoid this problem, make sure to enable the Pool Snapshot option for the archive server.

The Age of a File Version Exceeds Retention Thresholds

If your Delete Policy is excessive, file versions might exceed retention thresholds too quickly to be of use to your users. Typically, the file's archived versions are deleted automatically from the archive database as they exceed the Maximum Time to Keep Versions or the Maximum Number of Versions.

To avoid this problem, make sure your Delete Policy meets the needs of your users.

4.15 What's Next

After your deployment plan and resources meet the prerequisites and guidelines, you are ready to install Archive and Version Services. See the following:

- Chapter 5, "Installing and Configuring New Archive Servers," on page 47

Installing and Configuring New Archive Servers

5

This section describes how to install and configure Novell® Archive and Version Services 2.1 for NetWare®.

5.1 Before You Begin

Before installing Novell Open Enterprise Server NetWare and Novell Archive and Version Services 2.1 for NetWare:

1. Plan your archive server implementation.

   For information, see the following:

   - “Planning for Archive and Version Services” on page 25
   - “Coexistence and Migration Issues for Archive and Version Services” on page 115
   - “Security Considerations for Archive and Version Services” on page 119
2. Make sure that your planned archive server implementation satisfies the “Prerequisites and Guidelines” on page 37 for Novell Archive and Version Services.
3. Continue with Section 5.2, “Selecting a New Install Scenario,” on page 47.

5.2 Selecting a New Install Scenario

Depending on your planned implementation, continue with one of the following ways to set up an archive server:

Table 5-1 *New Install Scenarios*

Existing Archive	Single or	Refer To
None	Single server	Installing and Configuring an Archive Server (page 47)
None	Server cluster	Installing and Configuring an Archive Server Cluster (page 105)

5.3 Installing and Configuring an Archive Server

This section discusses the following tasks:

- Section 5.3.2, “Configuring Software RAID Devices for Your Archive Pool and Volume,” on page 49
- Section 5.3.3, “Creating an Archive Pool,” on page 50
- Section 5.3.4, “Creating an Archive Volume,” on page 51
- Section 5.3.5, “Configuring the ArkSQL Configuration File,” on page 52
- Section 5.3.6, “Installing and Configuring the ArkSQL Server,” on page 53
- Section 5.3.7, “Configuring Archive Server Information,” on page 56
- Section 5.3.8, “Archiving File Versions,” on page 58

5.3.1 Installing OES NetWare and ArkManager 2.1

After you plan your system and meet prerequisites and guidelines, you are ready to install the Archive and Version Services software module, ArkManager, on your server.

1. Install OES NetWare on your archive server, using the *Basic* install option.

   The Basic option installs ArkManager 2.1 on your server, but you must configure other key components after the install before you can run the program.

2. Confirm the install by looking for these elements in the `sys:\arkManager` directory.

   `arkConfig_sample_basic.xml`

   `arkConfig_sample_full.xml`

   `arkSQL_sample.cnf`

3. Confirm the install of the Archive Versioning plug-in for iManager, which you use as the management interface to control versioning jobs and to view the ArkManager job log.

   3a Launch a Web browser, then open it to the Novell iManager Login:

   `https://svr1.example.com/nps/iManager.html`

   Replace *svr1.example.com* with the actual IP address or DNS name of your archive server. The URL path is case sensitive.

   3b Log in as the administrator user (such as admin) to the Novell eDirectory™ tree that contains your archive server.

   3c The Archive Versioning role should be available in the Infrastructure category. For information about iManager, see Novell iManager 2.7 Administration Guide (http://www.novell.com/documentation/imanager27/index.html?page=/documentation/imanager27/imanager_admin_27/data/bsxrjzp.html#bsxrjzp).

4. Depending on your implementation plan, continue with one of the following:
   - Configuring Software RAID Devices for Your Archive Pool and Volume (page 49)
   - Creating an Archive Pool (page 50)

5.3.2 Configuring Software RAID Devices for Your Archive Pool and Volume

After you install OES NetWare, you can optionally create a software RAID device for the archive pool and volume to satisfy your availability needs. Use one of the following methods:

- Creating a Software RAID 1 or 5 Device (page 49)
- Creating a Software RAID 10 (page 49)

Creating a Software RAID 1 or 5 Device

To create a fault-tolerant solution for data, create a software RAID 1 or 5 device to use for your archive pool and volume. For information, see Managing Software RAID Devices in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

1. In iManager, expand the *Storage* role, then select *Software RAIDs*.
2. If it is not already selected, select the archive server.
3. Click *New* to open the Create a Software RAID dialog box.
4. Specify a name for the RAID.
5. Specify the type of RAID:
   - *RAID 1* (mirroring)
   - *RAID 5* (striping with parity)
6. (Conditional) For a RAID 5 device, specify the *Stripe Size*.

   The default size of 64 KB typically provides the best performance for devices with NSS volumes.
7. From the available devices, select the devices that you want to use.

   For a RAID 1 device, specify 2 to 4 devices. For a RAID 5 device, specify 3 to 14 devices.
8. Specify the amount of space to use for each segment.

   Each segment contributes equal amounts of space.
9. Click *OK*.
10. Continue with Section 5.3.3, “Creating an Archive Pool,” on page 50.

Creating a Software RAID 10

To provide maximum data fault tolerance, create a software RAID 10 (mirrored RAID 0 device) by creating a pool on a RAID 0 device, and then mirroring the pool’s partition.

1. Create 2 to 4 software RAID 0 (striping) devices. Repeat the following these steps to create each device:
   - **1a** In iManager, expand the *Storage* role, then select *Software RAIDs*.
   - **1b** If it is not already selected, select the archive server.
   - **1c** Click *New* to open the Create a Software RAID dialog box.
   - **1d** Specify a name for the RAID.
   - **1e** Specify the type of RAID as *RAID 0*.

1f Specify the *Stripe Size*.

The default size of 64 KB typically provides the best performance for devices with NSS volumes.

1g From the available devices, select the devices that you want to use.

Make sure that the segments in each of your RAID 0 devices have no devices in common; otherwise, you cannot mirror them later.

1h Specify the amount of space to use for each segment.

Each segment contributes equal amounts of space.

1i Click *OK*.

2 Create an archive pool on one of the RAID 0 devices. For information, see Section 5.3.3, "Creating an Archive Pool," on page 50.

3 Mirror the archive pool to create the RAID 10.

For information, see Mirroring an Existing Pool with NSSMU in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

3a At a server command prompt, enter

```
nssmu
```

3b In NSSMU, select *Partitions* from the NSSMU main menu.

3c Select the NSS partition for the archive pool.

3d Press F3 to create the RAID 1 device and mirror the partition.

3e From the available devices, select 1 to 3 of the RAID 0 devices you created in Step 1, then press Enter.

4 Continue with Creating an Archive Volume.

5.3.3 Creating an Archive Pool

On your OES NetWare archive server, you must create an NSS pool where you plan to store the archive database and archive data. For information, see Managing NSS Pools in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front) .

1 In iManager, expand the *Storage* role, then select *Pools*.

2 If it is not already selected, select the archive server.

3 Click *New* to open the New Pool wizard.

4 Specify a name for the pool. For example, *ark*.

5 Specify the devices to use.

If you created a software RAID device in Section 5.3.2, "Configuring Software RAID Devices for Your Archive Pool and Volume," on page 49, make sure to select that device for your pool.

6 Specify the amount of space to use.

7 Click *OK*.

8 Continue with Creating an Archive Volume.

5.3.4 Creating an Archive Volume

On your archive pool, create an NSS volume and directory where you plan to store the archive database and archive data. You should use the volume exclusively for the archive.

1 Create an NSS volume.

The following procedure describes how to create a nonencrypted NSS volume with iManager. For detailed information, see Creating and Configuring Unencrypted NSS Volumes with iManager (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/bqpdoh2.html#bqpdoh2) in the OES 2: *Novell Storage Services File System Administration Guide for OES* (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

If your implementation requires an encrypted NSS volume to store file versions from encrypted NSS data volumes, use the NSS Management Utility to create the encrypted volume. The iManager Storage plug-in does not provide an encryption option. For information, see Creating an Encrypted Volume (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/bq2y6nb.html#bq2y6nb) in the *OES 2: Novell Storage Services File System Administration Guide* (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/front.html#front).

1a In iManager, expand the *Storage* role, then select *Volumes*.

1b If it is not already selected, select the archive server.

1c Click *New* to open the New Volume wizard.

1d Configure the new volume.

- Specify a name for the volume. For example, *ark*.
- Select the NSS pool you created in Section 5.3.3, “Creating an Archive Pool,” on page 50, then select the *Allow the Volume to Grow to Pool Size* check box.
- Specify the desired attributes for the volume.

1e Click *Finish*.

2 (Optional, recommended) For easy backup, create a directory in the archive volume where you want to store the archive database and archive data.

For detailed information, see Creating a Directory (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/bs3pfti.html) in the *OES 2: File Systems Management Guide* (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/hn0r5fzo.html).

2a Open your Web browser to the Novell Remote Manager Login page on the archive server, and then log in with your administrator username and password. For example, enter

```
https://192.168.1.1:8009
```

Replace *192.168.1.1* with the actual IP address or DNS name of your archive server.

2b Click *Manage Server* > *Volumes*.

2c Click the *Properties* icon next to the archive volume.

2d Type the name of the directory, then click *Create Subdirectory*.

3 Continue with the next section, Configuring the ArkSQL Configuration File.

5.3.5 Configuring the ArkSQL Configuration File

The ArkSQL configuration file (`sys:\arkManager\arkSQL.cnf`) contains the configuration information for the Novell Archive and Version Services instance of MySQL on your MySQL server. ArkManager uses this information to access the archive database.

1 In a text editor, configure ArkSQL with the MySQL settings for the archive database.

1a Copy the contents of the `sys:\ArkManager\arkSQL_sample.cnf` file to the `sys:\arkManager\arkSQL.cnf` file.

Microsoft Windows automatically hides the `.cnf` extension. Ifyou use a Microsoft Windows computer to create or edit the arkSQL.`cnf` file, make sure file name is `arkSQL.cnf`, not `arkSQL.cnf.cnf`.

To view a copy of the `arkSQL_sample.cnf` file, see Section B.1, “Sample of the Database Configuration for ArkSQL.cnf,” on page 141.

1b Modify the `arkSQL.cnf` file to meet your needs. For example:

- Set the Data Directory path by modifying the following directives:

```
datadir = ark:/archive/

innodb_data_home_dir = ark:/archive/

innodb_log_group_home_dir = ark:/archive/

innodb_log_arch_dir = ark:/archive/
```

The data directory is the directory in your archive volume where the archive database resides. A directory location facilitates easy backup. Replace *ark* with the name of the archive volume. Replace *archive* with the name of the archive database directory.

- Set the Port Number by modifying the following directive:

```
port = 3306
```

Replace *3306* with the actual port number.

- Set the size of reserved storage space for the archive database by modifying the following directive:

```
innodb_data_file_path = ibdata1:400M:autoextend
```

Replace *400* with the amount of space (in MB) that you want to reserve initially for the archive database. The reserved size you specify must be at least 10 MB and must not exceed the volume size.

- (Optional) If your files are potentially very big, such as several hundred megabytes to gigabytes in size, consider increasing the time that queries wait to access a locked MySQL database. Set the `Lock Wait Timeout` variable by modifying the following directive:

  ```
  set-variable = innodb_lock_wait_timeout=50
  ```

  The default value is 50 seconds. Specify a longer wait time in seconds that is sufficient to avoid `Lock Request Timeout` errors. The more and bigger the files versioned, the longer the wait time should be.

2 Continue with the next section, Installing and Configuring the ArkSQL Server.

5.3.6 Installing and Configuring the ArkSQL Server

After you configure the ArkSQL configuration file, you are ready to install the ArkManager instance of MySQL on your archive server. Novell Archive and Version Services 2.1 uses a MySQL server to organize and host file versions in the archive database. MySQL is an open source, structured query language (SQL) database.

IMPORTANT: This guide refers to the Archive and Version Services instance of MySQL as ArkSQL.

For more information about managing MySQL, see the *OES 2: Novell MySQL for NetWare Administration Guide* (http://www.novell.com/documentation/oes2/web_mysql_nw/index.html?page=/documentation/oes2/web_mysql_nw/data/bookinfo.html#bookinfo) .

1 (Conditional) If MySQL is not installed on your archive server, install MySQL Server and configure an instance of MySQL for Novell Archive and Version Services. Otherwise, go to Step 2.

1a From the OES NetWare GUI, click the *NetWare post-install program*, then select *MySQL*.

The installation begins and opens the MySQL installation dialog box.

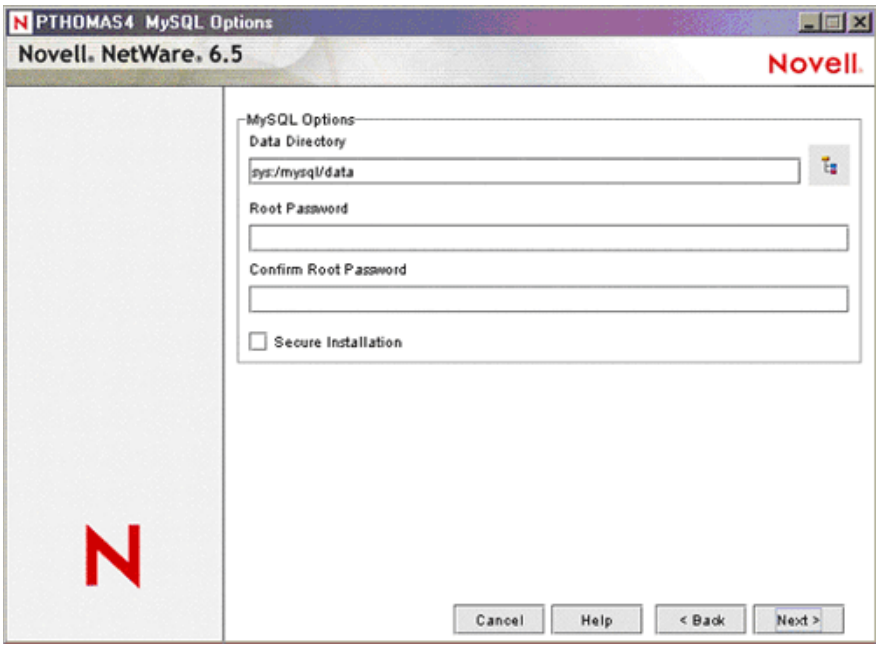

1b Complete the following information:

- **Data Directory:** The data directory is the directory in your archive volume where the archive database resides.A directory location facilitates easy backup.

  For example, type *ark:\archive*, where *ark* is the name of your archive volume and *archive* is the archive database directory. This is the same information you used in the `sys:\arkManager\arkSQL.cnf` file.

  IMPORTANT: Do not use the default location of `sys:\mysql\data`.

  Allow a minimum of 1 gigabyte of space on the archive volume for the data organization and management information related to ArkSQL.

- **Root Password:** The password for the root user, or superuser, who has access rights to perform any function.

  We strongly recommend that a password be assigned for the MySQL root user in a production environment. If you leave this field blank, anyone can connect as root without a password and be granted all privileges.

- **Port Number:** The port you want ArkManager to use for communications with ArkSQL. For example, type

  3306

  If other services use Port 3306, you can assign any free port, such as 3307 or 3308.

- **Secure Installation:** Select the *Secure Installation* check box to enable this option.

  IMPORTANT: We strongly recommend a secure installation for your production environment. If Secure Installation is enabled, it overrides the default installation by enforcing the requirement for a root password and restricting post-install access to the database to authorized users and only from the local host.

1c Click *Next*, then follow the on-screen instructions.

1d When the installation is complete, the MySQL service is running and the archive database is initialized.

1e If the ArkSQL instance is running at the intended port, shut it down. At the server console, enter

```
mysqladmin -p shutdown --port=value
```

Specify `-p` if the root user's password is set. Specify the `--port` only if the port is not 3306.

1f If the `ibdata1` file exists in the archive database directory, delete it.

For example, use a file manager to look for *`ark:`*`\archive\ibdata1`, where *`ark:`*`\archive` is the database directory you specified when you installed MySQL service. If it exists, delete it.

1g Restart ArkSQL service using `sys:\arkManager\arkSQL.cnf` file. At the server console, enter

```
mysqld_safe --defaults-file=sys:\arkManager\arkSQL.cnf
```

1h Go to Step 3.

2 (Conditional) If MySQL server is currently installed on your archive server, use MySQL commands to set up an instance of MySQL for Novell Archive and Version Services. Otherwise, go to Step 3.

2a If a MySQL instance is running at the intended port, shut it down. At the server console, enter

```
mysqladmin -p shutdown --port=value
```

Specify `-p` if the root user's password is set. Specify the `--port` if the port is not 3306.

2b To set the database directory path and port number and to initialize the database in that location: At the server console, enter

```
mysql_install_db --datadir=ark:\archive\ --port=3306
```

Replace *`ark`* with the name of the archive volume. Replace *`archive`* with the name of the archive database directory. Replace *`3306`* with the actual port number. These are the same values that you used in the `sys:\arkManager\arkSQL.cnf` file.

2c Start ArkSQL service using `sys:\arkManager\arkSQL.cnf` file. At the server console, enter

```
mysqld_safe --defaults-file=sys:\arkManager\arkSQL.cnf
```

2d To set a password for user "root": At the server console command prompt, enter

```
mysql -u root --port=3306
```

Replace *3306* with the actual port number.

2e At the MySQL prompt, enter

```
set password for root@'localhost'=password('your-password');
```

Replace *your-password* with the root user's actual password. (The ending semicolon is necessary.)

2f When the installation is complete, the MySQL service is running and the archive database is initialized on the archive volume.

3 Modify the archive server's `autoexec.ncf` file.

3a To start the ArkSQL server automatically whenever the OES NetWare server starts, add the following line in the `auotexec.ncf` file:

```
mysqld_safe --defaults-file=sys:\arkManager\arkSQL.cnf
```

3b Depending on the port number you are using for the ArkSQL instance of MySQL, do one of the following:

- If you are using Port 3306 for your ArkSQL server, comment out the following line by placing a pound sign (#) at the beginning of the line, as shown.

  ```
  #mysqld_safe --autoclose
  ```

- If you are using another port for your ArkSQL server, comment out the `mysqld_safe` command that matches your port number. For example:

  ```
  #mysqld_safe --port=3307 --autoclose
  ```

  Replace *3307* with your actual port number.

4 Continue with the next section, Configuring Archive Server Information.

5.3.7 Configuring Archive Server Information

After you install MySQL and configure ArkSQL, you are ready to configure the archive server information. Before ArkManager can run, the `sys:\arkManager\arkConfig.xml` file must contain the Novell eDirectory authentication information for ArkManager and the MySQL authentication information for the archive database. For information, see Section 3.3, "Understanding Archive Server Properties," on page 28.

The `arkConfig.xml` file and its parent directory are protected from unauthorized access by the file system trustees and trustee rights set for the file and directory. For information about the Novell trustee model for access control, see Directory and File Attributes for NSS Volumes or NetWare Traditional Volumes (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/bs3fih1.html) in the *OES 2: File Systems Management Guide* (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/hn0r5fzo.html).

The mandatory elements are:

- `<arkConfig>`

- <basic>
- <authentication> and its child elements:
  - <eDirectory> and its child elements: <user>, <password>, and <tree>
  - <database> and its child elements: <user>, <password>, and <portNumber>

The <displayLog/> tag is optional.

The passwords are read by ArkManager the first time it runs. It stores the information locally, then removes the passwords from the file. If you ever modify the passwords for the ArkManager administrator user or the MySQL database administrator user, you must add the information back to the arkConfig.xml file. For information, see Section 9.7, "Updating Passwords in ArkManager," on page 103.

References

During the configuration, refer to the following resources:

- For information about using XML, see Section 7.1, "Working with XML: An XML Primer," on page 77.
- For an example of the basic XML setup for the archive server information, see Section B.2, "Sample of a Basic Configuration for ArkConfig.xml," on page 144.
- For information about the basic ArkManager XML elements and their attributes, see Appendix A, "XML Elements and Attributes for ArkConfig," on page 121.

Procedure

1. In a text editor, copy the XML tags from the sys:\arkManager\arkConfig_sample_full.xml file to the sys:\arkManager\arkConfig.xml file.
2. In the arkConfig.xml file, modify the example settings for each of the mandatory elements with the actual values of your system.
3. Review your XML code to make sure that all required tags are present and that tags are properly formed.
4. Save the sys:\arkManager\arkConfig.xml file.
5. Confirm your ArkManager setup in a test environment before you create jobs for the archive server.
   - **5a** Make sure MySQL is running and at the correct port. It should start automatically when you boot your system. If it is not running, go to the server console and enter

     ```
     mysqld_safe --defaults-file=sys:\ArkManager\arkSQL.cnf
     ```

   - **5b** At the server console, enter

     ```
     arkstart
     ```

   - **5c** Verify that the ArkManager process starts.
   - **5d** At the server console, enter

```
arkstop
```

For information, see Section 9.2, “Stopping ArkManager,” on page 99.

6 Your archive server is ready to use. To manage the archive server settings in the future, see Chapter 9, “Managing the Archive Server,” on page 99.

7 Continue with the next section, Archiving File Versions.

5.3.8 Archiving File Versions

To archive file versions, you must define versioning jobs to run on the archive server. For information, see “Configuring Jobs in iManager” on page 59 or Chapter 7, “Configuring Jobs in ArkConfig,” on page 77.

To manage jobs, see Chapter 8, “Managing Jobs,” on page 85.

Configuring Jobs in iManager

6

To make recent file versions available to your users, you must set up one or more Novell® Archive and Version Services servers. This section discusses how to configure versioning jobs, using the Archive Versioning plug-in for Novell iManager.

6.1 Configuring the Archive Server Authentication Information

Before you run ArkManager for the first time, you must manually configure the mandatory Novell eDirectory™ authentication information for ArkManager and the MySQL database authentication information for ArkSQL in the `sys:\arkManager\arkConfig.xml` file.

The `sys:\arkManager\arkConfig_sample_basic.xml` file provides an example of the minimum set of XML elements to configure. To view an annotated version of the `arkConfig_sample_basic.xml` file, see Section B.2, “Sample of a Basic Configuration for ArkConfig.xml,” on page 144.

The following are mandatory elements:

- `<arkConfig>`
- `<basic>`
- `<authentication>` and its child elements:
  - `<eDirectory>` and its child elements: `<user>`, `<password>`, and `<tree>`
  - `<database>` and its child elements: `<user>`, `<password>`, and `<portNumber>`

The `<displayLog/>` tag is optional. For information, see “XML Elements and Attributes for ArkConfig” on page 121.

In the `arkConfig.xml` file, replace the example settings with your actual eDirectory and ArkSQL settings.

1. Copy the contents of the `sys:\arkManager\arkConfig_sample_basic.xml` file to the `sys:\arkManager\arkConfig.xml` file.
2. Modify the settings for each of the mandatory elements with the actual values of your system.

3 Save the `sys:\arkManager\arkConfig.xml` file.

4 Configure other arkConfig elements, using one or both of the following methods.

- Use the Web-based management interface in the Archive Versioning plug-in for iManager. For information, see Section 6.2, "Accessing the Archive Versioning Plug-In in iManager," on page 60.
- Use a text-editor to directly edit the `sys:\arkManager\arkConfig.xml` file. For information, see "Configuring Jobs in ArkConfig" on page 77.

6.2 Accessing the Archive Versioning Plug-In in iManager

1 Open your Web browser to the following URL:

`https://`*`svrname.example.com`*`/nps/iManager.html`

Replace *svrname.example.com* with the actual DNS name or IP address (for example, 192.168.1.1) of the server where iManager is running.

IMPORTANT: The URL path is case sensitive.

For information, see Accessing iManager (http://www.novell.com/documentation/imanager27/imanager_admin_27/index.html?page=/documentation/imanager27/imanager_admin_27/data/bsxrjzp.html) in the *Novell iManager 2.7 Administration Guide* (http://www.novell.com/documentation/imanager27/imanager_admin_27/index.html?page=/documentation/imanager27/imanager_admin_27/data/hk42s9ot.html#hk42s9ot).

2 On the iManager Login page, log in to the eDirectory tree where the archive server you want to manage resides.

3 In the left navigator, expand Archive Versioning to show its tasks.

4 Click *Archive Jobs* or *Archive Server Properties*, depending on what task you want to perform.

5 Select the archive server you want to manage in the tree where you are logged in to iManager.

For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.

6.3 Selecting an Archive Server to Manage

In the Archive Versioning tasks, you must select an archive server in order to activate the functions on the iManager pages. Use one of these methods to select an archive server in the tree where you are logged in to iManager:

- Type the Novell eDirectory server object name for the server you want to manage, then click *Apply*. For example:

  ```
  svr1.example
  ```

- Click the *Search* icon to open the eDirectory Object Selector. Browse or search the list to locate the server object of the server you want to manage, then click the server name.
- Click the *Object History* icon to select a server you have recently managed.

6.4 Configuring or Viewing Archive Server Properties

The Archive Server Properties page displays the basic management information that applies to all jobs controlled by the selected archive server. The general server properties specify the storage location where the archive data (file versions) reside and whether to display the log on the logger screen.

- Section 6.4.1, “Accessing the Server Properties Page,” on page 61
- Section 6.4.2, “Setting Archive Server Properties,” on page 61
- Section 6.4.3, “Viewing Authentication Properties,” on page 62

6.4.1 Accessing the Server Properties Page

To manage the archive server’s properties:

1. In iManager, expand *Archive Versioning*, select *Archive Server Properties*, then click *Properties*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

6.4.2 Setting Archive Server Properties

1. On the Server Properties page, set the following server properties for the archive server:

Property	Description
Volume Path	Specifies the path to the volume on the archive server where the file versions are stored for retrieval by users. The directory can be the root of the volume or any other directory in the volume. For example, on NetWare®, a Novell Storage Services™ volume path might be `ark:\archive`. If you are setting up a clustered solution for the archive server using Novell Cluster Services™, make sure to specify the volume and path to the virtual storage location.
Display Server Archive Log Entries on the Server's Console	By default, the archive server records error, warning, and normal messages for all of its jobs in the Archive and Version Services log. If *Display Archive Log Entries on the Server's Console* is selected, the archive server prints the messages to the server's logger screen in addition to recording them in the log. This is the default setting. If *Display Archive Log Entries on the Server's Console* is deselected, the archive server does not print messages to the logger screen, but the messages are recorded in the log. For information, see Section 8.5, "Viewing the Archive Log," on page 90.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.4.3 Viewing Authentication Properties

1 On the Server Properties page, view the authentication information for Novell eDirectory and MySQL that you configured in the `sys:\arkManager\arkConfig.xml` file. For information, see Section 6.1, "Configuring the Archive Server Authentication Information," on page 59.

Property	Description
eDirectory	**User Name** Specifies the Novell eDirectory common name of the administrator user who has the appropriate rights to the original data location and to the archive data location. For example, `admin.servercontext` The archive server administrator user must have supervisor rights to all servers being accessed by the selected archive server. The user must be in the same tree as the archive server and the source servers.
Database	Specifies authentication information about the MySQL database for the Archive and Version Services server. **User Name** Specifies the administrator user of the MySQL database. For example, root. **Port** The port used by the arkSQL instance of MySQL on the archive server. By default, Port 3306 is used, but you might have specified a different port when you set up arkSQL. If Port 3306 is used by other services, an alternate port can be used, such as 3307 or 3308.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.5 Configuring Default Job Settings

Default job settings are property settings that can optionally be used in any individual job on the server. Go to an individual job to configure it to use defaults.

6.5.1 Accessing the Default Job Settings Page

To manage the archive server’s default job settings:

1 In iManager, expand *Archive Versioning*, then click *Archive Server Properties*.

2 Select the archive server you want to manage, then wait for the page to refresh.

For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

3 Select the *Default Job Settings* tab.

6.5.2 Setting Default Job Information

1 On the Default Job Settings page, specify the following default job information property.

Property	Description
Copy All Files the First Time the Job Is Run	If *Copy All Files* is selected, the job saves versions of all eligible files to the archive database the first time the job is run. Thereafter, the job saves file versions only for files that changed. If *Copy All Files* is deselected, the job saves versions only of eligible files that were changed during the epoch, such as files that were added, modified, renamed, deleted, or where trustee information changed.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.5.3 Setting Default Source Server Information

1 On the Default Job Settings page, specify the following default source server information properties.

Property	Description
Server	The host name of the source server where the data to be versioned is located. For example, `servername.context`. The source server is in the same tree as the Archive and Version Services server. If you specify a new source server for an existing job, the archived data becomes associated with the source volume on the new source server. Typically, you change the source server only when you have renamed the it.
Snapshot Pool	The name of the destination pool where snapshots of the source volume are created and temporarily maintained at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, `pusers_s1` or `pdata_s1`. If no snapshot pool name is specified, or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.5.4 Setting Default Run Schedule Information

1 On the Default Settings page, pecify when to start the job and the frequency for running the job. To set the frequency, you must specify only one of three scheduling options:

Property	Description
No Run Schedule	Select *No Run Schedule* to indicate that there is no default schedule available for the server.
Every	Select *Every* to enable a *Scheduled Interval*. Specify the elapsed time between the beginning of versioning processes for a job. In the Units drop-down list, select seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours. If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe that the job skips some versioning intervals, you can increase the interval between versions, or you can reduce the amount of data to be versioned by setting Filter properties in the job's section of the `sys:\arkManager\arkConfig.xml` file.
Time	Select *Time*, specify the start time when the job's version process begins, then specify one or more days of the week to run the job. ◆ **Start Time:** Click the double-arrows to navigate to the time by hours or click the single-arrow to navigate by 15-minute increments. For example, 00:00 AM. ◆ **Days of the Week:** Select the *Days of the Week* check box to select all days for a daily job run, or select only the check boxes next to one or more days of the week you want the job to run. Choices include *Monday*, *Tuesday*, *Wednesday*, *Thursday*, *Friday*, *Saturday*, *Sunday*, and *All*.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.5.5 Setting Default Delete Policy Information

The delete policy determines the retention of file versions by age or by number of versions. If a delete policy is set, the job runs a process to delete file versions according to its own Delete Schedule. The delete process is not related to the job's Run Schedule, which determines when file versions are saved from the source volume. The job's delete policy runs if the job is in a Running, Scheduled, or Stopped state. The job's delete policy does not run if a job is in the Clean Up Users, Not Configured, or Deleted state.

IMPORTANT: The Delete Schedule operates separately from the Run Schedule.

1 On the Default Settings page, specify one of the following options:

Property	Description
No Delete Policy	Select this option if you want to retain file versions indefinitely.
Define Job Delete Policy	Select this option if you want to configure a default delete policy, then specify the schedule for deleting file versions, if they are eligible for deletion. Set the *Interval* and *Maximum Keep* settings.

2 If you selected *Defing Job Delete Policy*, complete the following information:

Property	Description
Interval	This setting represents the amount of time to wait from the time a delete-versions process ends until another delete-versions process begins. If a value is not specified, 24 hours is the default interval. How long the delete process takes depends on the number of files stored in the archive database. For example, suppose you set the *Delete Schedule Interval* to 1 hour. When you save a job, ArkManager starts the interval timer. After 1 hour elapses, the job kicks off its delete-versions process, resets the interval timer, and pauses the timer. When the delete process ends, the interval timer begins. The delete intervals repeat in this manner until the job is stopped.
Maximum Keep	Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the *Maximum Keep Versions* and *Maximum Keep Time* to zero values, the Delete Policy function does not run. ◆ **Time:** The maximum time that a file version is maintained in the archive. Specify the time in whole numbers; then in the *Units* drop-down list, select seconds, minutes, hours, or days. ◆ **Versions:** Specify the maximum number of versions of each file to keep in the archive. As the number of versions exceeds this integer value, the oldest version is deleted. ◆ **Keep Current Copy:** By default, the latest file version of current files remains in the database, even if the *Maximum Keep Time* elapses. If *Keep Current Copy* is selected, the archive keeps the latest file version as long as its source file is current on the source volume, even if the *Maximum Keep Time* elapses. After the user deletes the current source file, the deletion is noted when the versioning process runs. If the file version's age is within the *Maximum Keep Time*, the archive database retains the file version of the deleted file; otherwise, the file version is deleted. If *Keep Current Copy* is deselected, the archive deletes the file version when the *Maximum Keep Time* elapses.

3 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.5.6 Setting Default Filter Information

To specify the types of files and directories to version, set filters in the `sys:\arkManager\arkConfig.xml` file, using the Filter element and attributes as child elements of the Defaults element. For information, see Appendix A, "XML Elements and Attributes for ArkConfig," on page 121.

6.5.7 Applying Default Job Settings

Each job can optionally use none, one, or more of the default settings. Each job's usage of defaults is independent of other jobs' usage.

It is not necessary to stop jobs that use defaults while you modify, add, or remove default settings. Make your changes, then click *OK* or *Apply* to save them. The following table describes special circumstances for how the changes take effect:

Table 6-1 *How Changes to Default Job Settings Take Effect*

Change Made to Default Settings	State of the Job When You Save Default Job Settings	Additional Actions in Jobs that Use Defaults
Modifying an existing default Delete Policy	Running, Scheduled, or Stopped	If a delete-versions process is in progress, the run is completed with the old delete policy. The delete policy applies for the next delete-versions.
Modify any existing default setting except Delete Policy	Running	The job runs are completed with the old setting.
	Clean Up Users	The Clean Up Users run is stopped gracefully, and the job is placed in the Stopped state. The list might be only partially cleaned up.
	Not Configured	The job remains in the Not Configured state until you verify the individual job settings and start or schedule the job.
Add a default setting	Any	To take advantage of the new defaults, you must modify settings in individual jobs to use them.
Remove default settings	Running	The job run is completed with the old settings, then the job is placed in the Not Configured state. Go to the individual jobs to specify values for the missing settings.
	Scheduled or Stopped	Jobs are placed in the Not Configured state. Go to the individual jobs to specify values for the missing settings.
	Clean Up Users	The Clean Up Users run is stopped gracefully, and the job is placed in the Stopped state. The users list might be only partially cleaned up. The job is then placed in the Not Configured state. Go to the individual job to specify values for the missing settings.
	Not Configured	The job remains in the Not Configured state until you verify the individual job settings and start or schedule the job.

6.6 Configuring Job Properties

Use the following guidelines for setting job properties on the Create New Job page and Edit Job page.

6.6.1 Accessing the Job Settings Page

To manage the archive server's default job settings:

1. In iManager, expand *Archive Versioning*, then click *Archive Server Properties*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Job Settings* tab.

6.6.2 Setting Job Information

1. On the Job Settings page, configure the following job information properties:

Property	Description
Name	The administrator-specified unique job name. For example, svr1_users or svr1_data.
Copy All Files the First Time the Job Is Run	Initially, this field is set to the value specified in the Default Job settings. You can change it to the value you want to use for this job. If *Copy All Files* is selected, the job saves versions of all eligible files to the archive database the first time the job is run. Thereafter, the job saves file versions only for files that changed. If *Copy All Files* is deselected, the job saves versions only of eligible files that were changed during the epoch, such as files that were added, modified, renamed, deleted, or where trustee information changed.
Do Not Run	If *Do Not Run* is selected (enabled), the job enters the Stopped state when you save the job. You can start or schedule the job at any time thereafter. If *Do Not Run* is deselected (disabled), the job enters the Scheduled state when you save the job. However, if any setting is missing or invalid, the job enters the Not Configured state when you save the job, regardless of the *Do Not Run* setting. Edit the job to make corrections.

2. When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.6.3 Setting Source Server Information

1. On the Job Settings page, configure the following source server information properties:

Property	Description
Use Default Job's Source Server	If the default job defines a default source server, the *Use Default Job's Source Server* check box is selected. The default value is displayed in the field and you cannot modify the value. Deselect *Use Default Job's Source Server* to modify the values for that field. If the default job does not define a default source server, the *Use Default Job's Source Server* option is deselected and disabled.
Server	The host name of the source server where the data to be versioned is located. For example, `servername.context`. The source server is in the same eDirectory tree as the Archive and Version Services server. If you specify a new source server for an existing job, the archived data becomes associated with the source volume on the new source server. Typically, you change the source server only when you have renamed it.
Volume	The name of the source volume (or mount point) where the data to be versioned is located. For example, `users` or `data`.
Use Default Job's Snapshot Pool	If the default job defines a default snapshot pool, the *Use Default Job's Snapshot Pool* check box is selected. The default value is displayed in the field and you cannot modify the value. Deselect *Use Default Job's Snapshot Pool* to modify the values for that field. If the default job does not define a default snapshot pool, the *Use Default Job's Snapshot Pool* option is deselected and disabled.
Snapshot Pool	The name of the destination pool where snapshots of the source volume are created and temporarily maintained at the end of an epoch until point-in-time file versions can be saved to the archive database. For example, `pusers_s1` or `pdata_s1`. If no snapshot pool name is specified, or if the snapshot cannot be created for any reason, the versioning process copies files directly from the source volume. In this case, open files cannot be versioned.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.6.4 Setting Run Schedule Information

1 On the Job Settings page, specify when to start the job and the frequency for running the job. To set the frequency, you must specify only one of three scheduling options:

Property	Description
Use Default Job's Run Schedule	Select *Use Defaults* to use the *Run Schedule* settings specified on the Default Job Settings page. If there is no defined default schedule, this option is disabled.
Every	Select *Every* to enable *Scheduled Interval*. Specify the elapsed time between the beginning of versioning processes for a job. In the *Units* drop-down list, select seconds, minutes, hours, or days. For example, 45 seconds, 1 minute, 15 minutes, 2 hours, or 12 hours. If the versioning process exceeds the time specified as the interval, the overlapping scheduled job is skipped. No file versions are saved for skipped job runs. After the version process completes, the job runs at its next scheduled interval. If you observe that the job skips some versioning intervals, you can increase the interval between versions or reduce the amount of data to be versioned by setting Filter properties in the job's section of the `sys:\arkManager\arkConfig.xml` file.
Time	Select the *Time* field, then specify the start time when the job's version process begins, then specify one or more days of the week to run the job. ◆ **Start Time:** Click the double-arrows to navigate to the time by hours or click the single-arrow to navigate by 15-minute increments. For example, from 00:00 AM, click the double-arrow to move to 01:00 AM, then click the single arrow three times to reach 01:45 AM. ◆ **Days of the Week:** Select the *Days of the Week* check box to select all days for a daily job run, or select only the check boxes next to one or more days of the week you want the job to run. Choices include *Monday*, *Tuesday*, *Wednesday*, *Thursday*, *Friday*, *Saturday*, *Sunday*, and *All*.

2 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.6.5 Setting Delete Policy Information

The delete policy determines the retention of file versions by age or by number of versions. If a delete policy is set, the job runs a process to delete file versions according to its own Delete Schedule. The delete process is not related to the job's Run Schedule, which determines when file versions are saved from the source volume. The job's delete policy runs if the job is in a Running, Scheduled, or Stopped state. The job's delete policy does not run if a job is in the Clean Up Users, Not Configured, or Deleted state.

IMPORTANT: The Delete Schedule operates separately from the Run Schedule.

1 Use the Job Settings page to specify one of the following options:

Property	Description
No Delete Policy	Select this option if you want to retain file versions indefinitely.

Property	Description
Use Default Job's Delete Policy	Select *Use Defaults* to use the *Delete Policy* settings specified on the Default Job Settings page. If there is no defined default delete policy, this option is disabled and *Keep Latest Version of Current File* is selected by default.
Define Job Delete Policy	Specifies the schedule for deleting file versions, if they are eligible for deletion. Set the *Interval* and *Maximum Keep* settings.

2 If you selected *Define Job Delete Policy*, complete the following information:

Property	Description
Interval	This value represents the amount of time to wait from the time a delete-versions process ends until another delete-versions process begins. If a value is not specified, 24 hours is the default interval. How long the delete process takes depends on the number of files stored in the archive database. For example, suppose you set the Delete Schedule interval to 1 hour. When you save a job, ArkManager starts the interval timer. After 1 hour elapses, the job kicks off its delete-versions process, resets the interval timer, and pauses the timer. When the delete process ends, the interval timer begins. The delete intervals repeat in this manner until the job is stopped.
Maximum Keep	Specifies the maximum number of versions of each file to keep in the archive and how long to keep file versions. At least one of the values must be non-zero. If you set both the *Maximum Keep Versions* and *Maximum Keep Time* to zero values, the Delete Policy function does not run. ◆ **Time:** The maximum time that a file version is maintained in the archive. Specify the time in whole numbers; then in the Units drop-down list, select seconds, minutes, hours, or days. ◆ **Versions:** Specify the maximum number of versions of each file to keep in the archive. As the number of versions exceeds this integer value, the oldest version is deleted. ◆ **Keep Current Copy:** By default, the latest file version of current files remains in the database, even if the *Maximum Keep Time* elapses. If *Keep Current Copy* is selected, the archive keeps the latest file version as long as its source file is current on the source volume, even if the *Maximum Keep Time* elapses. After the user deletes the current source file, the deletion is noted when the versioning process runs. If the file version's age is within the *Maximum Keep Time*, the archive database retains the file version of the deleted file; otherwise, the file version is deleted. If *Keep Current Copy* is deselected, the archive deletes the file version when the *Maximum Keep Time* elapses.

3 When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

6.6.6 Setting Filter Information

To specify the types of files and directories to version, set filters in the `sys:\arkManager\arkConfig.xml` file, using the Filter element and attributes as child elements of the individual job's Jobs element. For information, see Appendix A, "XML Elements and Attributes for ArkConfig," on page 121.

6.7 Creating a Job

If your implementation plan calls for multiple archive servers, you must configure each with its own set of archive jobs.

Before you configure individual jobs, make sure you have configured the following information:

- **Archive Server Information:** Configure basic information for your archive server. For information, see Section 6.1, "Configuring the Archive Server Authentication Information," on page 59.
- **Default Job Settings:** Default settings can optionally be used by any job on the archive server. Each job can optionally use one or more of the defaults that are set; each job's usage is independent of other jobs.

  To use the defaults, you must set default job settings before the job runs. (In *Roles and Tasks*, expand the *Archive Versioning* role, select *Archive Server Properties*, the select *Default Job Settings*.) For information, see Section 6.5, "Configuring Default Job Settings," on page 63.

To create a job:

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Jobs* tab.
4. Click *New* to open the Create New Job page.
5. Specify the property settings to use when archiving file versions for a specified volume.

   For information about property settings, see Section 6.6, "Configuring Job Properties," on page 67.

   You must at least specify a job property or use the default job property for all required fields before the job can run. Required fields are marked with an asterisk (*).
6. When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

   The job is saved only if all required fields are completed and all settings are valid.
7. If necessary, correct any missing or invalid information, then click *Apply*.

   For information about annotations for missing or invalid information, see Section 6.8.1, "Correcting Missing or Invalid Information," on page 73.

6.8 Editing a Job's Properties

After you configure a job, you might occasionally need to modify its properties, such as to change the run schedule or delete policy. If you modify the server's default job settings, it can affect jobs that use those defaults.

6.8.1 Correcting Missing or Invalid Information

For existing jobs, the Job Properties page marks fields that are missing information or contain invalid information.

Required Fields

You must specify a job property or use the default job property for a required field before the job can run. Required fields are marked with an asterisk (*).

Not Configured Fields

The job might report a field as Not Configured if the setting is missing or invalid, or if the default job setting the field uses are missing or invalid. While a job is in the Not Configured state, the job and its delete policy do not run.

Settings that need attention are marked with three asterisks (***) on a red background. Verify the job settings and make any desired changes. If necessary, go to the *Archive Server Properties* > *Default Job Settings* to edit default job settings.

For example, fields marked as Not Configured might have missing or invalid data under the following circumstances:

- If you define a job directly in the `sys:\arkManager\arkConfig.xml` file and a required setting is missing or invalid, the job is placed in a Not Configured state. (If you define jobs in the Archive Versioning plug-in to iManager, the interface does not allow you to save a job with invalid or missing values.)
- If you remove settings in the *Default Job Settings*, jobs that use the defaults are affected.
- If a source volume is not mounted, it appears that the setting is invalid.
- If the source server is down, it appears that the setting is invalid.
- If you previously removed a job's definition directly in the `sys:\arkManager\arkConfig.xml` file instead of deleting the job with iManager, the job's data remains in the database, but it has no job definition.

Defaults

To use the default job properties, you must set default job settings before the job runs. In *Roles and Tasks*, click *Archive Server Properties*, then select *Default Job Settings* to configure default properties.

6.8.2 Modifying Job Settings

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Jobs* tab.
4. (Optional) Select *Archive Server Properties*, then select *Default Settings* to verify or set default job settings you want to use for the job.

   For information, see Section 6.5, "Configuring Default Job Settings," on page 63.
5. On the *Jobs Report*, select a *Job* check box, then click *Edit* to open the Edit Job page.
6. Specify settings for all required fields and for any optional settings you want to use.

   For information, see Section 6.6, "Configuring Job Properties," on page 67.

   You must specify a valid job property, or use the default job property, for all required fields before the job can run. Required fields are marked with an asterisk (*).
7. Correct any invalid settings.

   Invalid information is marked with three asterisks (***) on a red background. For information, see Section 6.8.1, "Correcting Missing or Invalid Information," on page 73.
8. When you are done, click *Apply* to save your changes, or click *Cancel* at any time to discard them.

   The job is saved only if all required fields are completed and all settings are valid.

6.8.3 Applying Modified Job Settings

It is not necessary to stop jobs while you modify their settings. If all required and optional settings are valid:

- The job is rescheduled forward from the save time.
- If a delete-versions job is in progress, the run is completed with the old delete policy. A modified delete policy is applied when the delete job runs again.
- You can expect the following additional actions, depending on the Do Not Run setting and the state of the job when you save your changes:

Table 6-2 *How Changes to Job Settings Take Effect*

Do Not Run Setting	State of the Job When You Save Edits	Additional Actions
Enabled (selected)	Running	The running job is completed with the old settings, then the job enters the Stopped state.
	Scheduled, Stopped, or Not Configured	The job enters the Stopped state.
	Clean Up Users	An in-progress cleanup of the job's users stops gracefully when you begin the edit, and the job enters the Stopped state. When you save changes, the job remains in the Stopped state.

Do Not Run Setting	State of the Job When You Save Edits	Additional Actions
Disabled (deselected)	Running	The running job is completed with the old settings, then it enters the Scheduled state.
	Scheduled, Stopped, or Not Configured	The job enters the Scheduled state.
	Clean Up Users	An in-progress cleanup of the job's users stops gracefully when you begin the edit, and the job enters the Stopped state. When you save changes, the job enters the Scheduled state.

6.9 What's Next

To manage Archive and Version services jobs, see "Managing Jobs" on page 85.

Configuring Jobs in ArkConfig 7

This section discusses how to configure jobs for Novell® Archive and Version Services for NetWare®, using an XML (Extensible Markup Language) configuration file. Use the `sys:\arkManager\arkConfig.xml` file to control the versioning processes for the archive server it is on. Configure the basic settings, the job defaults, and one or more individual jobs for volumes with data that you want to version.

7.1 Working with XML: An XML Primer

An XML document is a plain text file that contains hierarchically structured data. You can edit the file in a text editor if you are using ASCII characters. For non-ASCII characters, you must convert the XML document to UTF-8 encoding.

An XML document describes the data it contains using elements and attributes. An element identifies what the data is and an attribute defines metadata about the data. The relative placement of elements within the XML structure matters because an element can take on different meaning, based on where it is located in the XML structure.

For information, see the following:

7.1.1 Elements

An element uses a set of markup tags to delimit and define each piece of data in a document. The tag set consists of an start tag and an end tag. For example:

```
<tagname>data</tagname>
```

An element consists of information from the start tag to the end tag and everything in between.

If data contains the less than (<), greater than (>), or ampersand (&) characters, you must enclose the data with the CData tag.

```
<![CDATA[data]]>
```

For example:

```
<tagname><![CDATA[1&2]]></tagname>
```

7.1.2 Attributes

Elements can be annotated with any number of unique attributes. Attributes appear as name/value pairs separated by an equal sign (=) and must appear in double quotes or single quotes. You attach attributes in the start tag, but not to the end tag. For example:

```
<tagname attribute_name="value">data</tagname>
```

7.1.3 Hierarchical Relationships between XML Elements

Elements have a hierarchical structure that defines the relationships between parent and child elements. Every XML document has exactly one top-level element, known as the root element. The root element is mandatory, even if it has no content. All other elements are its children.

Some elements appear only once in a document, while others can appear multiple times. The child elements of the root element can be a parent to multiple elements. An element can be a child of different parent elements.

7.1.4 Element Content

Element content is the information between the two tags of an element, such as no content, parsed character data (PCData), and child elements. If a tag has no element content, it is called an empty element. The table below shows examples of some common XML tag constructs.

Element Content	Example of Tag Construct	Description
No content	`<tagname></tagname>`	An empty element
	`<tagname/>`	A abbreviated form of an empty element
Parsed character data (PCData)	`<tagname>text</tagname>`	An element with data
	`<tagname attribute_name="value">`	An element with data and with one or more attributes that describe the data
	`text`	
	`</tagname>`	

Element Content	Example of Tag Construct	Description
One or more child elements	<tagname> <child1_tag>text</child1_tag> <child2_tag></child2_tag> </tagname>	An element with two child elements: an element with PCData and an empty element

7.1.5 Rules for Well-Formed XML Documents

As you configure the XML file, keep in mind the following rules for well-formed XML documents:

- Every XML document has only one root element.
- Every start tag must have a matching end tag. The exception is the abbreviated version of an empty element (<tagname/>).
- Tags cannot overlap; every element must be properly nested.
- Element names and attribute names are case sensitive.
- XML keeps any white space in your text.

7.1.6 Additional Information

For more information about using XML, consult an XML programming textbook or search the Internet for an XML tutorial.

7.2 Configuring the ArkConfig.xml File

Use the sys:\arkManager\arkConfig.xml file to control which data resources to version, to set a schedule for versioning, and to determine the lifetime for versions, such as by age or by numbers of versions.

The XML root element in the arkConfig.xml file is the arkConfig (<arkConfig>) element. Within the root element, you will have at least three child elements:

- **Basic Element**

  Defines the file system authentication information and the archive volume location on your archive server. Specify the physical location of the archive volume for single-server solutions. Specify the virtual location of the archive volume for clustered-server solutions.

  The Basic element also contains any other non-job-specific settings such as whether entries to the log should also be displayed on the server's logger screen. All child elements in the Basic element are mandatory except for <displayLog/>.

- **Defaults Element**

  Defines the settings for the server context to use, the location of data to version, and the frequency that versioning occurs. Use this element to specify the parameters to use for version jobs if they are not specified in the individual Job elements.

- **Job Element**

  Defines the settings for versioning of a particular volume or of filtered data within the volume.

XML is hierarchical in nature; information is structured like a tree, with parent-child relationships. The figure below illustrates the XML hierarchy for an `arkConfig.xml` file. See Appendix A, "XML Elements and Attributes for ArkConfig," on page 121 for a detailed discussion of each of these elements.

You can use the sample `sys:\arkManager\arkConfig_sample_full.xml` file as a guide to configure the settings for your archive server; do not attempt to use the sample as is. For an annotated version of this sample XML file, see Section B.3, "Sample of a Full Configuration for ArkConfig.xml," on page 146.

Figure 7-1 *Hierarchy of Elements in the arkConfig.xml File*

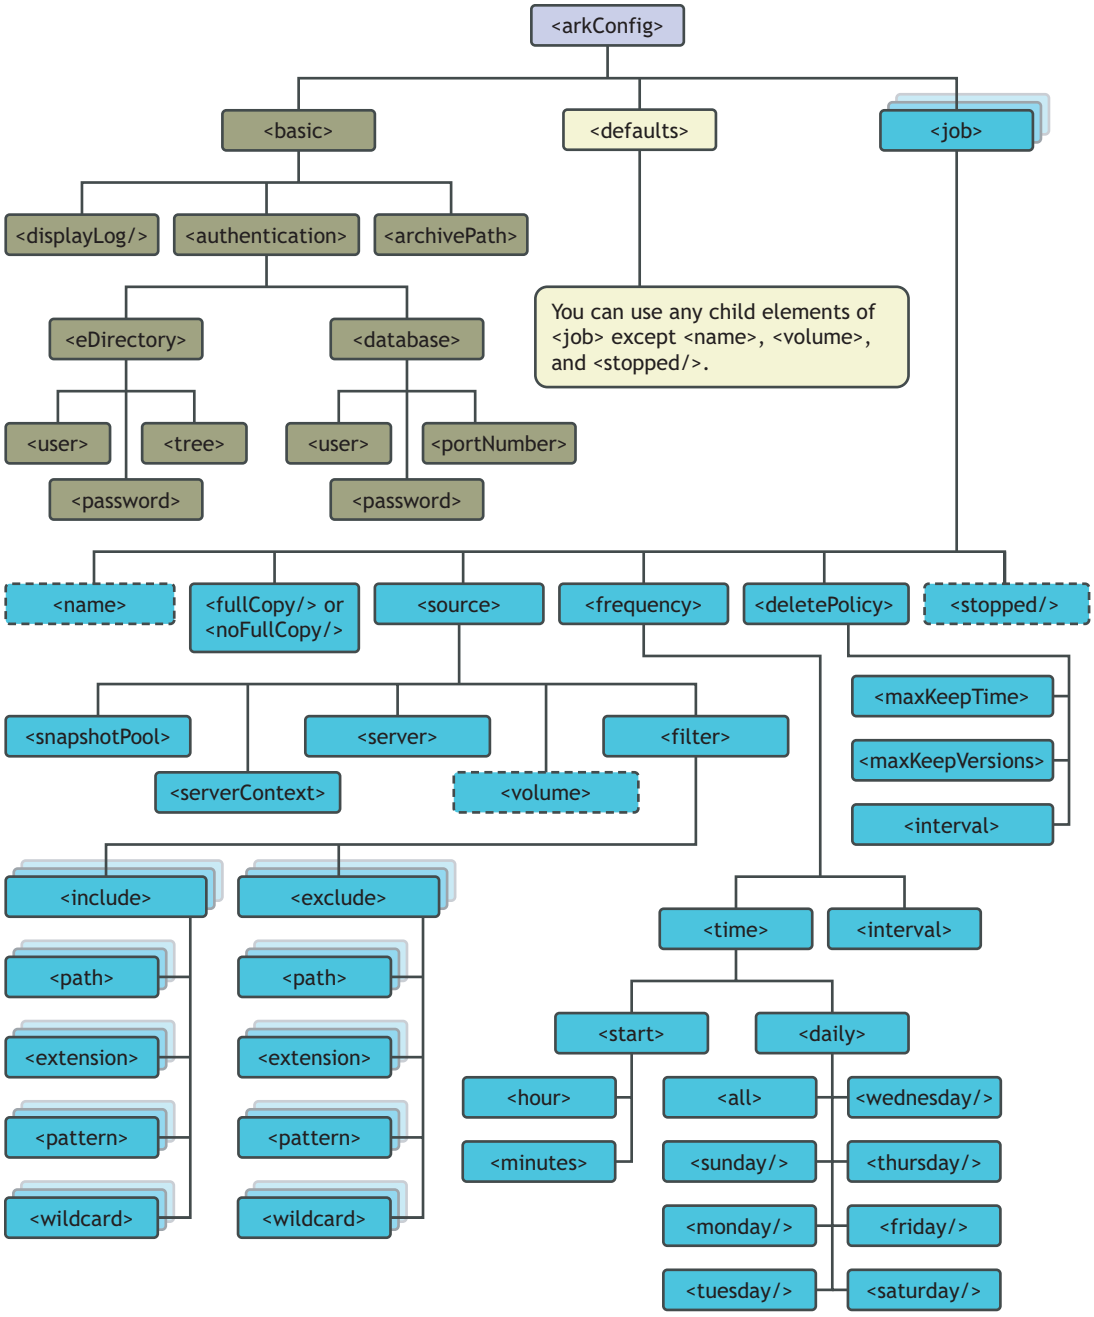

1 Before you begin, make sure your versioning plan satisfies the following guidelines:

- Section 4.9, “Job Guidelines,” on page 42
- Section 4.10, “Schedule Guidelines,” on page 43
- “Security Considerations for Archive and Version Services” on page 119

2 Stop the archive server. At the server console, enter

```
arkstop
```

For information, see Section 9.2, "Stopping ArkManager," on page 99.

3 In a text editor, open the `sys:\arkManager\arkConfig.xml` file.

4 (Conditional) Configure the archive server information, if necessary, by modifying data in the Basic element.

You initially configured this information when you set up your server. You need to modify this information only if you have modified the password for the admin user of Novell eDirectory or if you have modified the root user name, password, or port number for your MySQL server. For information, see Section 5.3.7, "Configuring Archive Server Information," on page 56.

5 Configure the *Default Job Settings* to apply for job parameters when individual jobs on this archive server do not specify a value to use. To do this, modify information in the Defaults element.

When you first set up Defaults, you can copy the Defaults element from the `arkConfig_sample_full` file, and then modify it as needed. For information, see Section 3.4, "Understanding Job Properties," on page 29.

6 For each source volume, define a single Job element for the selected archive server.

When you first set up a job, you can copy the Job element from the `sys:\arkManager\arkConfig_sample_full.xml` file, and then modify it as needed.

6a Provide a unique name for the job.

6b Modify the job settings. For information, see Section 3.4, "Understanding Job Properties," on page 29.

6c Make sure to set the job's Start time at some time in the future so that you have time to correct any job settings before the versioning begins.

7 Save your changes.

8 Make sure the XML code you entered in `sys:\arkManager\arkConfig.xml` file is well-formed. Look for missing tags, illegal characters, and missing characters.

Some Web browsers can be used to validate XML code, but they cannot verify that the ArkManager settings are correct. For information validating XML with your browser, consult your browser's Help.

Correct any mistakes to create a well-formed XML document, then save your changes.

9 Start ArkManager. At the server console, enter

```
arkstart
```

10 Verify that ArkManager recognizes the jobs.

10a View a list of jobs with the Archive Versioning plug-in for Novell iManager. For information, see Section 8.1, "Viewing a Jobs Report," on page 85.

10b If the job is in the list, continue with Step 11.

If the job is not in the list, you probably have an XML error. Continue with Step 10c.

10c Stop ArkManager. At the server console, enter

```
arkstop
```

10d In a text editor, open the `sys:\arkManager\arkConfig.xml` file.

10e Review the XML code for the job definition.

10f Make the necessary corrections, then save the file.

10g Start ArkManager. At the server console, enter

```
arkstart
```

10h Repeat this process until the missing job appears in the Jobs list.

11 Verify that the job settings are correct.

11a View a job's details with the Archive Versioning plug-in for Novell iManager. For information, see Section 8.2, "Viewing a Job's Details," on page 87.

11b If the job details are correct, continue with Step 12.

If the job details are not correct, you probably have an XML error. Continue with Step 11c.

11c Stop ArkManager. At the server console, enter

```
arkstop
```

11d In a text editor, open the `sys:\arkManager\arkConfig.xml` file.

11e Review the XML code for the job definition.

11f Make the necessary corrections, then save the file.

11g Start ArkManager. At the server console, enter

```
arkstart
```

11h Repeat this process until the job details are correct.

12 Start the job now or schedule the job. For information, see Section 8.3, "Starting or Scheduling a Job," on page 88.

7.3 What's Next

To manage Archive and Version services jobs, see Chapter 8, "Managing Jobs," on page 85.

Managing Jobs

8

After creating jobs for Novell® Archive and Version Services 2.1 for NetWare®, manage them with the Archive Versioning plug-in for Novell iManager.

This section discusses the following tasks:

8.1 Viewing a Jobs Report

The jobs report lists all current jobs on a selected archive server and reports their status and schedule.

8.1.1 Generating the Jobs Report

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.
3. Select the *Jobs* tab.
4. View a report of all its current Archive and Version Services jobs.

   Each entry in the table represents a single job.

8.1.2 Sorting the Jobs Report by Column

All columns are sortable in ascending and descending order. Click a heading link to sort the jobs by that column. Click the link a second time to sort the jobs in reverse order. A sort icon next to the heading indicates which column is being used as the sort key and the sort order.

8.1.3 Setting the Refresh Rate for the Jobs Report

You can control the refresh frequency of the jobs report. The setting persists only while you are viewing the report or while the browser session is current.

Select *Refresh*, then select the desired frequency from the drop-down list:

Off
Immediately
Every 5 seconds
Every 15 seconds
Every 30 seconds [Default]
Every 60 seconds
Every 5 minutes
Every 15 minutes

8.1.4 Understanding the Report Content

The following table describes information in the Jobs Report:

Table 8-1 *Description of Information in the Jobs Report*

Property	Description
Jobs	Select a *Job* check box to select the job you want to manage.
Name	The administrator-specified name of the job. Typical names might provide information about the source server and source volume. For example: *server_volume* or *srv1_vol1*. Click the job's *Name* link to edit the job.
Status	 orts the current operational state of the job. Click the *Status* icon to view the job's details.  Scheduled: The job is scheduled to save file versions; it is not currently running.  Running: The job is actively saving file versions to the archive database.  Stopped: The job's archived data is available to users, but the job is disabled so that no new versions are being captured. Cleanup: The Clean Up Users job is in progress. For information, see Section 8.7, "Cleaning Up the Job's User List," on page 94. Not Configured: One or more of the job's settings are missing or invalid, or the default job settings the job uses are missing or invalid. While a job is in the Not Configured state, the job and its Delete policy do not run.

Property	Description
Status (cont'd)	Settings that need attention are marked with three asterisks (***) on a red background. Click the job's *Name* link to view or edit the job properties. Go to the Default Job Settings page (Archive Versioning > Archive Server Properties > Default Job Settings) to edit default job settings. Verify the job settings and make any necessary changes. For example, fields marked as Not Configured might have missing or invalid data under the following circumstances: ◆ If you define a job directly in the `sys:\arkManager\arkConfig.xml` file and a required setting is missing or invalid, the job is placed in a Not Configured state. (If you define jobs in the Archive Versioning plug-in to iManager, the interface does not allow you to save a job with invalid or missing values.) ◆ If you remove settings in the Default Job Settings, jobs that use the defaults are affected. ◆ If a source volume is not mounted, it appears that the setting is invalid. ◆ If the source server is down, it appears that the setting is invalid. ◆ If you previously removed a job's definition directly in the `sys:\arkManager\arkConfig.xml` file instead of deleting the job with iManager, the job's data remains in the database, but it has no job definition.
Next Start Date	The next scheduled date and time that the job is scheduled to begin a versioning process. The format is mm/dd/yy hh:mm xm. For example: 4/25/05 11:45 PM. If the job is in a Stopped state, the *Next Start Date* might have a date and time that is previous to the current time. Whenever the job starts again, you can specify whether to run the job immediately (*Run Now* or *Run Now with Copy All* actions) or to schedule the job to run later (*Run Later* or *Run Later with Copy All* actions).
Source Server - Source Volume	The server name and volume name of the volume that the job archives. A source volume can have only one archive job defined for it.

8.2 Viewing a Job's Details

The Job Details page shows the effective settings for a job, which are a combination of the job's individual properties and the default settings the job uses.

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Jobs* tab.
4. View a report of all its current Archive and Version Services jobs.

   Each entry in the table represents a single job.
5. Use one of the following methods to view the details of an Archive Job:
   - Select a *Job* check box, then click *Actions* > *Details*. Select only one job at a time.
   - Click the *Status* icon for the job.

   The *Job Details* box displays the following information about the selected job.

Property	Description
Name	The administrator-specified name of the job.
Status	Reports the current operational state of the job.  *Scheduled*: The job is scheduled to save file versions; it is not currently running.  *Running*: The job is actively saving file versions to the archive database.  *Stopped*: The job's archived data is available to users, but the job is disabled so that no new versions are being captured. *Cleanup*: The Clean Up Users job is in progress. For information, see Cleaning Up a Job's Users List. *Not Configured*: One or more of the job's settings are missing or invalid, or the default job settings the job uses are missing or invalid. While a job is in the *Not Configured* state, the job and its Delete policy do not run.
Source Server	The eDirectory Server object name of the source server where the data to be versioned is located. The source server is in the same tree as the Archive and Version Services server.
Source Volume	The name of the volume where the data to be versioned for this job is located.
Snapshot Pool	The name of the destination pool where snapshots of the source volume are created and temporarily maintained until point-in-time file versions can be saved to the archive database.
Last Start Date	The date and time that the job started the last time it was successfully run and completed. If the job is in the Running state, this time stamp is not updated until the job completes successfully.
Next Start Date	The next scheduled date and time that the job is scheduled to begin a versioning process. The format is mm/dd/yy hh:mm xm. For example: 04/25/05 11:45 PM. If the job is in a Stopped state, the *Next Start Date* might have a date and time that is previous to the current time. Whenever the job starts again, you can specify whether to run the job immediately (*Run Now* or *Run Now with Copy All* actions) or to schedule the job to run later (*Run Later* or *Run Later with Copy All* actions).
Schedule	Reports the interval or start times for the selected job. • **Scheduled Interval:** The elapsed time between the beginning of versioning processes for a job. • **Scheduled Start Time:** The time of day when the job's version process begins and the days of the week the job runs.

6 When you are done, click *Close*.

8.3 Starting or Scheduling a Job

1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.

2 Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

3 Select the *Jobs* tab.

4 View a report of all its current Archive and Version Services jobs.

   Each entry in the table represents a single job.

5 In the *Jobs Report*, select one or more *Job* check boxes, click *Actions*, then select one of the following actions:

Property	Description
Run Now	Starts the selected jobs immediately.
Run Later	Schedules the selected jobs to run at their individual next scheduled start times.
Run Now with Copy All	Starts the selected jobs immediately and copies all files specified by the individual selected jobs from their source volumes to the archive volume on a one-time, special version occurrence.
Run Later with Copy All	Schedules the selected jobs to run at their individual next scheduled start times. For that run only, it copies all files specified by the individual selected jobs from their source volumes to the archive volume on a one-time, special version occurrence.

8.4 Stopping a Job

Stop a job whenever you want to temporarily pause versioning or to cease versioning a particular source volume.

1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.

2 Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

3 Select the *Jobs* tab.

4 View a report of all its current Archive and Version Services jobs.

   Each entry in the table represents a single job.

5 In the *Jobs Report*, select one or more *Job* check boxes, then click *Actions* > *Stop*.

Stopping a job stops the selected jobs gracefully and disables them from running again until you start or schedule each job. The Stopped setting persists through a server reboot. For information, see Section 8.3, “Starting or Scheduling a Job,” on page 88.

Stopping a job does not stop the delete policy from running. To stop the delete policy, modify the job’s *Delete Policy* setting to *No Delete Policy*. For information, see Section 6.8, “Editing a Job’s Properties,” on page 73.

You can also use the `<stopped/>` tag in the *Job* element for the particular job in the `sys:\arkManager\arkConfig.xml` file.

8.5 Viewing the Archive Log

ArkManager issues messages as it performs tasks for each job. The Archive Log displays error, warning, and normal messages reported by all jobs on a selected archive server. You can filter log entries to view only those messages of interest (such as by severity type or by job), to set the number of messages to view per page, and to specify the date of the messages you want to jump to first in the log.

This section discusses the following topics:

- Understanding Log Entries
- Viewing Messages in the Archive Log
- Viewing Messages on the Server Logger Screen

8.5.1 Understanding Log Entries

The Archive Log contains all messages for selected jobs on a selected archive server. By default, the log displays messages for all severity types and for all jobs, beginning with the current date and time. Each log entry displays the following information:

Table 8-2 *Description of Parameters in the Archive Log*

Property	Description
Job Name	The administrator-specified unique name of the job.
Severity	Specifies the alert level of the job's message. Messages have three levels of severity: ◆ **Error:** Reports a failure to perform a job task. For example, the versioning service might not be able to save a version of a file from the source volume to the archive volume because of a broken connection. ◆ **Warning:** Reports a non-critical error in performing a job. ◆ **Normal:** Reports the status of normal job tasks such as starting, copying, and completing.
Date	The date and time stamp of the message. The format is mm/dd/yy hh:mm xm. For example: 04/25/05 11:45 PM.
Message	The text message issued by the job.

8.5.2 Viewing Messages in the Archive Log

- "Viewing a Log of All Jobs" on page 90
- "Viewing a Log of Selected Jobs" on page 91
- "Navigating the Archive Log" on page 91
- "Updating the Log Message Entries" on page 91
- "Filtering the Log Message Entries" on page 91

Viewing a Log of All Jobs

1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.

2 Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

3 Select the *Log* tab.

4 View the Archive and Version Services log, which contains all messages for all jobs on the selected server. You can apply a filter for message types or time periods.

Viewing a Log of Selected Jobs

1 In iManager, expand *Archive Versioning*, then click *Archive Jobs*.

2 Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, “Selecting an Archive Server to Manage,” on page 61.

3 Select the *Jobs* tab.

4 In a Jobs report, select one or more *Job* check boxes, then click *Actions* > *View Log*.

5 View the Archive and Version Services log.

   The log opens with a filter applied so that the log displays only messages for the selected jobs.

Navigating the Archive Log

By default, the log displays 10 messages per screen. Messages appear from newest (top and left) to oldest (bottom and right). Use the browser scroll bar to view all entries in the page. Use the arrows at the bottom of the page to move from page to page.

- Use the left-arrow to navigate to the previous page of newer messages.
- Use the double left-arrow to go directly to the newest messages.
- Use the right-arrow to navigate to the next page of older messages.
- Use the double right-arrow to go directly to the oldest messages.

Updating the Log Message Entries

To update the log you are viewing to include the most recent messages, click *Archive Job*, then click the *Log* tab. You can also click the *Jobs* tab, then click the *Log* tab.

Filtering the Log Message Entries

- **Filter:** You can view all available log entries, or use a filter to view only those messages you want to see. The log filter allows you to filter messages by severity types and by one or more jobs so that only the messages you want to view are displayed.

  By default, the log displays 10 messages per page and begins with the newest messages. You can use the filter to set preferences for the number of log entries to display per page and the date and time of the entries you want to see first in the log.

  To set the filter, click Filter. For information, see Section 8.6, “Filtering Messages in the Archive Log,” on page 92.

- **Clear Filter:** Resets the current filter settings to the default log display values, using the current date and time as the starting point of the log, and it displays 10 entries per page.

  To reset the filter to the default *Archive Log* values, click *Clear Filter*.

8.5.3 Viewing Messages on the Server Logger Screen

In addition to viewing the log in iManager, you can display the log on the server Logger screen. To enable this option, include the `<displayLog/>` element in the job's configuration in the `sys:\arkManager\arkConfig.xml` file. For information, see "Configuring Jobs in ArkConfig" on page 77.

This is an empty element. It has no child elements and defines no data values. Either the shortened version or the long version of this element is valid. For example:

```
<displayLog/>
```

or

```
<displayLog></displayLog>
```

8.6 Filtering Messages in the Archive Log

In the iManager Archive Log page, click *Filter* to filter messages so that the log displays only those log entries that you want to see. You can also set preferences for how many log entries to view per page and for the date and time of messages you want to jump to first. You can set one or more of the filter types each time you activate the filter process. Click *Cancel* at any time to back out of the filtering process.

The Filter page allows the following:

- Section 8.6.1, "Filtering Log Entries by Severity Type," on page 92
- Section 8.6.2, "Filtering Log Entries by Job," on page 93
- Section 8.6.3, "Setting the Display Details," on page 93
- Section 8.6.4, "Setting the Date and Time Range," on page 93
- Section 8.6.5, "Resetting the Log Filters," on page 94

8.6.1 Filtering Log Entries by Severity Type

Set filters so that the log displays only the messages that match one or more *Severity Types* that you select. Deselected types are not displayed.

1. On the Archive Log page, click *Filter*.
2. On the Filter Archive Log page, select one or more *Severity Type* check boxes next to the message types you want to display in the log.
3. (Optional) Set any other filter options on the page.
4. Click *OK* to save the settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.2 Filtering Log Entries by Job

The archive log displays recent messages for all jobs on the selected server. An ArkManager job appears in the list if there are messages that occurred that could not be allocated to a particular job. A filtered log displays messages for only the jobs you select; it filters out all jobs that are deselected.

1. On the Archive Log page, click *Filter*.
2. Select one or more *Job* check boxes next to the jobs you want to display in the log.
3. (Optional) Set any other filter options on the page.
4. Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.3 Setting the Display Details

1. On the Archive Log page, click *Filter*.
2. Specify the number of log entries to display per page.

   The default is 10 messages.
3. Specify one of the following:
   - Select *Display Log Messages* to display message text in the *Messages* column.
   - Deselect *Display Log Messages* if you want to omit the *Messages* column in the archive log.
4. (Optional) Set any other filter options on the page.
5. Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.4 Setting the Date and Time Range

The log automatically displays all job messages with the newest messages first. You can specify the date and time of the first message to display. The log jumps directly to the log entries from that date and time instead of to the newest messages. All messages are still available by using the arrow keys at the bottom of the screen to navigate to newer and older messages than the date you set.

To view the most recent log entries:

1. Click *View Newest Log Entries*.

To specify a particular date and time for entries to jump to first:

1. On the Archive Log page, click *Filter*.
2. To specify a date, click the *Calendar* icon, then select a date by using the arrows to navigate to the correct month and year.
3. To specify a time, click the arrows to navigate by hours (double arrows) or by 15-minute increments (single arrows) to the time.
4. (Optional) Set any other filter options on the page.
5. Click *OK* to save settings, or click *Cancel* at any time to back out of the filter setting process.

8.6.5 Resetting the Log Filters

The filter settings are not persistent. They are active only until you use one of the following methods to reset them:

- Use the *Filter* to specify new *Filter* settings.
- Click *Clear Filter* on the Archive Log page to return to default settings.
- Log out of iManager.
- Close the iManager browser window.

8.7 Cleaning Up the Job's User List

For each job, ArkManager maintains information about its users centrally instead of storing duplicate information among the file versions. The user list grows cumulatively; it is not automatically synchronized to remove users from the list who no longer have associations with current file versions in the archive. To clear out obsolete user information, run the Clean Up Users task about once a month, depending on the job's schedule, the job's Delete policy, and the turnover of users in your workplace.

During the cleanup, jobs are in the Clean Up Users state. While Clean Up Users is running, the job itself does not run and the job's Delete policy does not run. Afterwards, the jobs are automatically placed in the Stopped state and the jobs' Delete policies are resumed.

If a cleanup job is interrupted, the job stops gracefully and enters the Stopped state. The user list might be only partially cleaned up. For example, the job is interrupted if you delete, edit, start, or schedule the job, or if you shut down or restart ArkManager. If you are editing a job, when you save your edits, the job continues in the Stopped state if you enabled the *Do Not Run* setting, or it enters the Scheduled state if you disabled the *Do Not Run* setting.

To clean up the user list:

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Jobs* tab.
4. In the Jobs report, select one or more jobs where you want to clean up the user list, then select *Actions* > *Stop*.

   Wait for the jobs to stop gracefully and report a Stopped status.
5. Select one or more jobs, then select *Actions* > *Clean Up Users*.

   When the cleanup for a job is complete, the job enters the Stopped state.
6. When cleanup is done, restart the stopped jobs.

   For information, see Section 8.3, "Starting or Scheduling a Job," on page 88.

8.8 Deleting File Versions

Users can use the NSS File Versions Utility from their Windows workstations to select a file and delete all versions of the file from the archive database.

Administrators can delete file versions for a variety of reasons, such as the following:

- To clear the database of unnecessary data that was versioned prior to applying job filters
- To delete file versions without compromising the ArkSQL data structure
- To clear the database of file versions related to a job you plan to delete (See also Section 8.9, "Deleting a Job," on page 95.)
- To clear the database of file versions to free space in the archive volume

Deleting a Specified File or Directory Version Set

Use the NSS File Versions Utility or the NetStorage Archive function to locate and delete a selected file version set or a selected directory version and its contents from the archive database. For information, see *OES 2: Novell Archive and Version Services 2.1 User Guide* (http://www.novell.com/documentation/oes2/bkup_arc_user_lx_nw/index.html?page=/documentation/oes2/bkup_arc_user_lx_nw/data/front.html#front).

Deleting Older File Versions

Modify the *Maximum Keep Time*, *Maximum Keep Versions*, or *Keep Current Copy* settings in the job's *Delete Policy* settings to delete only the older file versions in the archive database. For information, see Section 6.8, "Editing a Job's Properties," on page 73.

Deleting All File Versions for a Job

Delete and purge the job to remove all of its file versions from the archive database. For information, see the following:

- Section 8.9, "Deleting a Job," on page 95
- Section 8.11, "Salvaging a Deleted Job," on page 97

8.9 Deleting a Job

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Jobs* tab.
4. In the Jobs Report, select one or more *Job* check boxes, click *Delete*, then confirm the deletion.

   If the job is running when it is deleted, ArkManager stops the job gracefully, then deletes the job. A partial data set exists for that run, up to the point where the job was stopped.

The deleted job is retained in the archive database in the Deleted job state, but it no longer appears in the list of current jobs in the Jobs Report. Its job name can be reused for a current job.

While a job is in the Deleted state:

- The deleted job does not run.
- Users cannot access the deleted job's file versions in the archive database.
- The deleted job's Delete Policy is not enforced.

A deleted job's properties and archived file versions remain in the archive database until you purge or salvage the job. For information, see the following:

- Section 8.10, "Viewing a Deleted Jobs Report," on page 96
- Section 8.11, "Salvaging a Deleted Job," on page 97
- Section 8.12, "Purging a Deleted Job," on page 98

8.10 Viewing a Deleted Jobs Report

A deleted job's properties and archived file versions remain in the archive database until you purge the job. While a job is in the Deleted state:

- The deleted job does not run.
- Users cannot access the deleted job's file versions in the archive database.
- The deleted job's Delete Policy is not enforced.

To view a list of deleted jobs:

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Deleted Jobs* tab.
4. View a report of all deleted jobs on the selected server.
5. Each entry in the table represents a single job. If multiple deleted jobs have the same job name, they are distinguishable by their deleted times, source servers, and source volumes.

   All columns are sortable in ascending and descending order. Click a heading link to sort the jobs by that column. Click the link a second time to sort the jobs in reverse order. A sort icon next to the heading indicates which column is being used as the sort key and the sort order.

Property	Description
Jobs	Select a *Job* check box to select the job you want to manage.
Name	The administrator-specified name of the job. Typical names might provide information about the source server and source volume. For example: server_volume or srv1_vol1. Click the job's *Name* link to edit the job.
Source Server - Source Volume	The server name and volume name of the volume that the job archived.
Date	The date and time when the deletion process was completed successfully.

6. When you are done, click *Close*.

8.11 Salvaging a Deleted Job

You can salvage a deleted job to restore the job and to allow users to access its archived file versions. You can salvage only one job at a time. You cannot salvage a deleted job if the source volume it archived is currently the source volume for an existing job on the archive versioning server.

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Deleted Jobs* tab.
4. In the Deleted Jobs Report, select the *Job* check box next to the job you want to salvage, then click *Salvage*.

   If multiple deleted jobs have the same job name, they are distinguishable by their deleted dates, source volumes, and source servers.
5. Specify a name for the salvaged job that is unique among the non-deleted jobs on the archive server.
6. Click *OK*.

   Salvaging a deleted job restores the job. The following table describes the state of the job after it is salvaged, given the job's state when it was deleted.

State of the Job When Deleted	Delete Action	State of the Job After Salvage
Running	The job stops gracefully, enters the Stopped state, then it is deleted. A partial data set exists for the job.	Scheduled The job's next run time is calculated forward from the time the job is salvaged.
	The job's Delete policy is suspended.	The job's Delete policy runs immediately.
Stopped	The job is deleted.	Stopped
	The job's Delete policy is suspended.	The job remains in the Stopped state until you start or schedule it. The job's Delete policy runs immediately.
Scheduled	The job is deleted.	Scheduled
	The job's Delete policy is suspended.	The job's next run time is calculated forward from the time the job is salvaged. The job's Delete policy runs immediately.
Not Configured	The job is deleted.	Not Configured
	The job's Delete policy, which does not run in the Not Configured state, continues to be suspended.	The job cannot be started or scheduled until you provide valid settings. The job's Delete policy does not run until you start or schedule the job.

State of the Job When Deleted	Delete Action	State of the Job After Salvage
Clean Up Users	The Clean Up Users job stops gracefully, enters the Stopped state, then it is deleted. The user list might be only partially cleaned up. The job's Delete policy, which does not run in the Clean Up Users state, is now suspended.	Stopped The job remains in the Stopped state until you start or schedule it. The job's Delete policy runs immediately.

8.12 Purging a Deleted Job

Purging a deleted job removes the job properties and all archived data for the job from the archive database. Purged jobs cannot be recovered.

1. In iManager, expand *Archive Versioning*, then click *Archive Jobs*.
2. Select the archive server you want to manage, then wait for the page to refresh.

   For information, see Section 6.3, "Selecting an Archive Server to Manage," on page 61.
3. Select the *Deleted Jobs* tab.
4. In the Deleted Jobs Report, select the *Job* check box next to the job you want to purge, then click *Purge*.

   If multiple deleted jobs have the same job name, they are distinguishable by their deleted dates, source volumes, and source servers.
5. Click *OK* to confirm the purge, or click *Cancel* to back out of the process.

Managing the Archive Server

9

This section discusses how to manage your Novell® Archive and Version Services 2.1 for NetWare® server.

9.1 Starting ArkManager

Starting ArkManager on a Server

To start ArkManager, enter the following command at the server console command prompt:

```
arkstart
```

The ArkSQL instance of MySQL starts automatically whenever you boot the server. If you need to start ArkSQL manually, see Section 9.3, “Starting the MySQL Server and ArkSQL,” on page 100.

If you want ArkManager and ArkSQL to start automatically on server reboot, add the `arkstart` command in the `autoexec.ncf` file.

Starting ArkManager and ArkSQL in a Cluster

For a clustered server solution, start ArkManager and ArkSQL by starting the cluster resources. For information, see the *OES 2: Novell Cluster Services 1.8.4 for NetWare Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

9.2 Stopping ArkManager

Stopping ArkManager on a Server

To stop ArkManager, enter the following command at the server console command prompt:

```
arkstop
```

After typing `arkstop`, make sure ArkManager is shutting down cleanly before running it again. You can check this in two ways:

- If *Display Log* is enabled for the server (that is, the `<displayLog/>` tag is placed in the `arkconfig.xml` file), go to the server's Logger screen and wait until the `ArkManager shutdown is completed` message appears.
- On the archive server's console screen, use the `java -show` command to check that no ArkManager Java thread is running.

It might take a while for ArkManager to shut down because it needs to wait for its current file archiving process to finish. Under normal circumstances, ArkManager should eventually shut down cleanly.

If you need to stop the ArkSQL instance of MySQL, see Section 9.4, "Stopping the MySQL Server and ArkSQL," on page 101.

Killing the ArkManager Process If It Hangs

If you suspect the shutdown process is taking an unreasonably long time, and it might be hanging, use the following procedure to get ArkManager's thread ID and kill the ArkManager process manually.

1. At the server console command prompt, enter

   ```
   java -show
   ```

   This returns a list of current process IDs.
2. Get ArkManager's process ID from the list.
3. At the server console command prompt, enter

   ```
   java -killxx
   ```

   where *xx* is ArkManager's process ID.

Stopping ArkManager and ArkSQL in a Cluster

For a clustered server solution, stop ArkManager and ArkSQL by stopping the cluster resources. For information, see the *OES 2: Novell Cluster Services 1.8.4 for NetWare Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

9.3 Starting the MySQL Server and ArkSQL

When you installed and configured the MySQL instance for ArkManager, you set up commands in the `autoexec.ncf` file to automatically start ArkSQL whenever you boot the system.

Manually Starting ArkSQL

If you need to manually start ArkSQL, at the server console, enter

```
mysqld_safe --defaults-file=sys:\ArkManager\arkSQL.cnf
```

Starting ArkSQL in a Cluster

For a clustered server solution, start ArkSQL by starting the cluster resources. For information, see the *OES 2: Novell Cluster Services 1.8.4 for NetWare Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

9.4 Stopping the MySQL Server and ArkSQL

Stopping the MySQL Server and ArkSQL on a Server

At the server console, enter

```
mysqladmin -p shutdown --port=value
```

Specify `-p` if the root user's password is set. Specify the `--port` if the port is not 3306.

Make sure that the MySQL Database Server screen has closed. If `mysqld_safe` command was not issued with the `--autoclose` option (a typical situation for ArkSQL), go to the MySQL Database Server screen and press any key to close the screen.

Stopping ArkSQL in a Cluster

For a clustered server solution, stop ArkSQL by stopping the cluster resources. For information, see the *OES 2: Novell Cluster Services 1.8.4 for NetWare Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

9.5 Modifying the ArkSQL Settings

If you need to modify the MySQL root username, password, or port number, you must modify the values in MySQL, in the `sys:\arkManager\arkSQL.cnf` file, and in the `sys:\arkManager\arkConfig.xml` file.

1. Stop ArkManager.

   For information, see Section 9.2, "Stopping ArkManager," on page 99.

2. If the ArkSQL instance is running at the intended port, shut it down. At the server console, enter

   ```
   mysqladmin -p shutdown --port=value
   ```

   Specify `-p` if the root user's password is set. Specify the `--port` if the port is not 3306.

3 In a text editor, configure ArkSQL with new settings, as necessary.

3a Set the port number by modifying the following directive:

```
port = 3306
```

Replace *3306* with the actual port number.

3b Save your changes.

3c (Conditional) If this is an archive server cluster, copy the revised `sys:\arkManager\arkSQL.cnf` file to each server in the cluster.

4 Update the archive server information in the `sys:\arkManager\arkConfig.xml` file. For information, see Section 9.6, "Modifying the Archive Server Information," on page 102.

5 To start MySQL, at the server console, enter

```
mysqld_safe --defaults-file=sys:\arkManager\arkSQL.cnf
```

6 Start ArkManager by entering `arkstart` at the server console.

9.6 Modifying the Archive Server Information

You must update the archive server information in the Basic elements of the `sys:\arkManager\arkConfig.xml` file if you modify any of the following information for the archive server:

- Novell eDirectory™ username or password
- MySQL root username, password, or port number

If you change the administrator user's password on the network or if you change the root user's password in MySQL, you must also update the passwords that arkManager uses when it starts. ArkManager cannot restart until you pass it the correct passwords in the `arkConfig.xml` file.

For information, see Section 3.3, "Understanding Archive Server Properties," on page 28 and Section B.2, "Sample of a Basic Configuration for ArkConfig.xml," on page 144.

1 Stop ArkManager.

For information, see Section 9.2, "Stopping ArkManager," on page 99.

2 Stop the MySQL instance for ArkSQL.

For information, see Section 9.4, "Stopping the MySQL Server and ArkSQL," on page 101.

3 In a text editor, open the `sys:\arkManager\arkConfig.xml`.

4 Modify the settings, as necessary, for each of the mandatory elements with the new values of your system.

5 Review your XML code to make sure that all required tags are present and that tags are well-formed.

6 Save the `sys:\arkManager\arkConfig.xml` file.

7 Start ArkManager.

For information, see Section 9.1, "Starting ArkManager," on page 99.

9.7 Updating Passwords in ArkManager

The administrator passwords for Novell eDirectory™ and MySQL are initially stored in the `sys:\arkManager\arkConfig.xml` file in clear text. The first time you run ArkManager, it reads the password, encrypts it, then stores the encrypted password in a local store. It removes the password from the `arkConfig.xml` file, then saves the file. Meanwhile, the `arkConfig.xml` file is protected from unauthorized access by the file system trustees and trustee rights that apply to the `sys:\arkManager` directory and the `arkConfig.xml` file itself. For information about the Novell trustee model for secure access, see Directory and File Attributes for NSS Volumes or NetWare Traditional Volumes (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/bs3fih1.html) in the *OES 2: File Systems Management Guide* (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/hn0r5fzo.html).

ArkManager refers to the encrypted passwords in the local store whenever you use the `arkstart` command. If you change the administrator user's password on the network or the root user's password for MySQL, you must also update the related passwords that arkManager uses when it starts. ArkManager cannot restart until you pass it the new passwords.

1. Stop ArkManager.

   For information, see Section 9.2, "Stopping ArkManager," on page 99.
2. Edit the `sys:\arkManager\arkConfig.xml` file to add the correct passwords in the `<password>` elements in the `<eDirectory>` section or the `<database>` section, as needed.
3. Start ArkManager.

   For information, see Section 9.1, "Starting ArkManager," on page 99.

9.8 Backing Up the Archive Database

You should periodically back up the archive database. The frequency depends on the critical nature of the versions you store in an archive database.

WARNING: To ensure data integrity, shut down ArkManager before and during the backup. For information, see Section 9.2, "Stopping ArkManager," on page 99.

For information, see Chapter 16.9: Backing Up and Recovering an InnoDB Database (http://dev.mysql.com/doc/mysql/en/Backing_up.html) in the *MySQL Reference Manual* (http://dev.mysql.com/doc/mysql/en/index.html).

9.9 Backing Up the Archive Data

You should periodically back up the archive data. The frequency depends on the critical nature of the versions you archive.

WARNING: To ensure data integrity, shut down ArkManager and ArkSQL before and during the backup. For information, see Section 9.2, "Stopping ArkManager," on page 99.

To back up the archive versions for a particular job, copy the directories to a different volume. We recommend that the destination volume where you store the backup copy be on a different drive and pool than your archive volume, although it is not mandatory. You can also use your standard backup tools and procedures to make the directories part of your scheduled backup.

The general data path is

```
ark:\archive\arkDataxxxxxx\
```

where *ark* is the archive volume, *archive* is the archive data directory, and *xxxxxx* is a 6-digit random number assigned to the job.

If the archive path is in a directory in the volume, simply copy all files in the directory.

9.10 Recovering the Archive Database

For information about recovering InnoDB databases such as ArkSQL, see Chapter 16.9: Backing Up and Recovering an InnoDB Database (http://dev.mysql.com/doc/mysql/en/Backing_up.html) in the *MySQL Reference Manual* (http://dev.mysql.com/doc/mysql/en/index.html).

9.11 Enabling UTF-8 Encoding Support for Clients

UTF-8 encoding support for clients makes it possible to access multiple language files from an NCP™ client without changing the code page on either the NCP client or the NetWare server. ArkManager uses UTF-8 encoding for filenames. If the language code pages of the NCP client and server are different, and if you are using an older version of the Novell Client™, the server conveys an incorrect filename to the archive database. In turn, the archive server presents unfamiliar filenames to the user who is trying to restore file versions.

To prevent the code page problem, upgrade your NCP clients to use the Novell Client in this release of Open Enterprise Server, which provides UTF-8 support, then enable the UTF-8 service. CIFS clients have included UTF-8 encoding support since NetWare 6.0 SP 2 and NetWare 6.5.

To enable the UTF-8 service:

1. Install or upgrade to the newest Novell Client in OES NetWare.
2. Log in to the Novell Client.
3. To open the Novell Client Configuration window, right-click the *Novell Client* icon in the system taskbar and then click *Novell Client Properties*.
4. Click *Advanced Settings*.
5. In *Parameter Groups*, select *Use UTF-8 Encoding and NCPs*.
6. In *Setting*, set the value to *On* to enable this service.
7. Click *OK* to apply the setting.
8. When prompted, reboot the client machine.

Installing and Configuring an Archive Server Cluster

10

You can configure the archive server in an active/passive cluster configuration, using Novell Cluster Services™ 1.8.5 for OES NetWare.

Before you attempt to implement this solution, familiarize yourself with how Cluster Services works. For information, see the *NW65 SP8: Novell Cluster Services 1.8.5 Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

This section discusses the following tasks:

10.1 Installing OES NetWare and ArkManager 2.1

After you plan your system and meet prerequisites and guidelines, you are ready to install the Archive and Version Services software module, ArkManager, on your server.

Perform the following procedure for each server you plan to include in the archive server cluster:

1. Install OES NetWare on your archive server, using the *Basic* install option.

   The Basic option installs ArkManager 2.1 on your server, but you must configure other key components after the install before you can run the program.

2. The configuration files and sample configuration files are located in the `sys:\arkManager` directory. Confirm the install by looking for the following elements:

   ```
   arkConfig_sample_basic.xml

   arkConfig_sample_full.xml

   arkSQL_sample.cnf
   ```

3 Use the Archive Versioning plug-in for iManager as the management interface to control versioning jobs and to view the ArkManager job log.

3a Launch a Web browser, then open it to the Novell iManager Login:

```
https://svr1.example.com/nps/iManager.html
```

Replace *svr1.example.com* with the actual IP address or DNS name of your archive server. The URL path is case sensitive.

3b Log in as the administrator user (such as admin) to the Novell eDirectory tree that contains your archive server.

3c The *Archive Versioning* role is available in the *Infrastructure* category. For information about iManager, see Novell iManager 2.7 Administration Guide (http://www.novell.com/documentation/imanager27/imanager_admin_27/index.html?page=/documentation/imanager27/imanager_admin_27/data/hk42s9ot.html#hk42s9ot).

4 Continue with the next section, Making Devices Sharable for Clustering.

10.2 Making Devices Sharable for Clustering

After you install OES NetWare, make your storage devices sharable for clustering.

1 In iManager, expand the *Storage* role, then click *Devices*.

2 Select the primary archive server for the cluster as the server you want to manage.

3 Select the devices you plan to use for the archive pool and volume, then mark the devices as *Sharable*.

Make sure the devices you choose are located in a configuration that is available to every server in the planned cluster. For more information, see Sharing Devices for NSS Pools (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/agtgjgq.html) in the *OES 2: NSS File System Administration Guide* (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

4 Depending on your implementation plan, continue with one of the following:

- Section 10.3, “Configuring Sharable Software RAID Devices for Your Archive Pool and Volume,” on page 106
- Section 10.4, “Installing Novell Cluster Services 1.7 for OES NetWare,” on page 108

10.3 Configuring Sharable Software RAID Devices for Your Archive Pool and Volume

After you make devices sharable for clustering, you can optionally create a software RAID devices for the archive pool and volume to satisfy your availability needs. Use one of the following methods:

- Creating a Sharable Software RAID 1 or 5 Device (page 107)
- Creating a Sharable Software RAID 10 (page 107)

10.3.1 Creating a Sharable Software RAID 1 or 5 Device

To create a data fault-tolerant solution, create a software RAID 1 or 5 device to use for your archive pool and volume. For information, see Creating a Software RAID Device with iManager (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/ajhvbk8.html) in the *OES 2: NSS File System Administration Guide* (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html) .

1. In iManager, expand the *Storage* role, then click *Software RAIDs*.
2. If it is not already selected, select the primary archive server in the cluster.
3. Click *New* to open the Create a Software RAID dialog box.
4. Specify a name for the RAID.
5. Specify the type of RAID:
   - *RAID 1* (mirroring)
   - *RAID 5* (striping with parity)
6. (Conditional) For a RAID 5 device, specify the *Stripe Size*.

   The default size of 64 KB typically provides the best performance for devices with NSS volumes.
7. From the available devices, select the devices that you made sharable in Section 10.2, “Making Devices Sharable for Clustering,” on page 106.

   For a RAID 1 device, specify 2 to 4 devices. For a RAID 5 device, specify 3 to 14 devices.
8. Specify the amount of space to use for each segment.

   Each segment contributes equal amounts of space.
9. Click *OK*.
10. Continue with Section 10.4, “Installing Novell Cluster Services 1.7 for OES NetWare,” on page 108.

10.3.2 Creating a Sharable Software RAID 10

To provide maximum data fault tolerance, create a software RAID 10 (mirrored RAID 0 device) pool by creating a pool on a RAID 0 device, and then mirroring the pool.

1. Create 2 to 4 RAID 0 (striping) devices. Repeat the following steps to create each device:
   - **1a** In iManager, expand the *Storage* role, then click *Software RAIDs*.
   - **1b** If it is not already selected, select the primary archive server in the cluster.
   - **1c** Click *New* to open the Create a Software RAID dialog box.
   - **1d** Specify a name for the RAID.
   - **1e** Specify the type of RAID as *RAID 0*.
   - **1f** Specify the *Stripe Size*.

     The default size of 64 KB typically provides the best performance for devices with NSS volumes.
   - **1g** From the available devices, select the devices that you made sharable in Section 10.2, “Making Devices Sharable for Clustering,” on page 106.

Make sure that the segments in each of your RAID 0 devices have no devices in common; otherwise, you cannot mirror them later.

1h Specify the amount of space to use for each segment.

Each segment contributes equal amounts of space.

1i Click *OK*.

2 Install Novell Cluster Services 1.7 on each of the servers. For information, see Section 10.4, "Installing Novell Cluster Services 1.7 for OES NetWare," on page 108.

3 Create a cluster-enabled pool on one of the RAID 0 devices. For information, see Section 10.5, "Creating a Sharable Archive Pool," on page 109.

4 Mirror the cluster-enabled pool to create the RAID 10.

For detailed information, see Mirroring an Existing Pool with NSSMU (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/ajhvbk8.html) in the *OES 2: NSS File System Administration Guide* (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html).

4a At a server console prompt, enter

```
nssmu
```

4b In NSSMU, select Partitions from the NSSMU main menu.

4c Select the NSS partition for the archive pool.

This is the cluster-enabled pool you create in Section 10.5, "Creating a Sharable Archive Pool," on page 109.

4d Press F3 to create the RAID 1 device and mirror the partition.

4e From the available devices, select 1 to 3 of the RAID 0 devices you created in Step 1, then press Enter.

5 Continue with Section 10.6, "Creating a Shared Archive Volume," on page 109.

10.4 Installing Novell Cluster Services 1.7 for OES NetWare

1 Install Novell Cluster Services 1.7 for OES NetWare for each server you plan to include in the archive server cluster.

For information, see the *NW65 SP8: Novell Cluster Services 1.8.5 Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

2 After Cluster Services is installed successfully on each server, continue with the next section, Creating a Sharable Archive Pool.

10.5 Creating a Sharable Archive Pool

Create a cluster-enabled, sharable pool. For more information, see Creating a Pool (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/bqpd5ps.html) in the *OES 2: NSS File System Administration Guide* (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html).

1. In iManager, expand the *Storage* role, then click *Pools*.
2. If it is not already selected, select the primary server in the cluster.
3. Click *New* to open the New Pool wizard.
4. Name the pool.
5. For each of the shared devices, specify the space to use in the pool.

   If you select sharable software RAID devices, the entire space is used in the pool.
6. Select *Cluster Enable on Creation* to enable this option.
7. Select *Activate on Creation* to enable this option.
8. To complete the cluster information, specify the following *Shared-Pool Clustering Parameters*:
   - **Virtual Server Name:** The name assigned by NetWare to the virtual server that represents the shared pool in the cluster.
   - **CIFS Virtual Server Name:** The name assigned by NetWare to the virtual server for handling CIFS (Common Internet File System) requests. This is the name of the server as it appears in a Windows system.
   - **IP Address:** The IP address that you want to assign the virtual server.

     To specify an IP address, tab between the different entries; no dot is required in the fields. For example, if the IP address is 192.168.1.1, type

     192 168 1 1
   - **Advertising Protocols:** Protocols that give users native file access to data: NCP, CIFS, and AFP.

     Specify one or more advertising protocols by selecting the check boxes of the protocols you want to enable for data requests to this shared pool.
9. Click *Finish*.
10. Continue with the next section, Creating a Shared Archive Volume.

10.6 Creating a Shared Archive Volume

On the cluster-enabled archive pool, create an NSS volume and directory where you plan to store the archive database and archive data. You should use the volume exclusively for the archive.

1. Configure an NSS volume.

   The following procedure describes how to create a nonencrypted NSS volume with iManager. For detailed information, see Creating and Configuring Unencrypted NSS Volumes with iManager (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/bqpdoh2.html#bqpdoh2) in the OES 2: *Novell Storage*

Services File System Administration Guide for OES (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

If your implementation requires an encrypted NSS volume to store file versions from encrypted NSS data volumes, use the NSS Management Utility to create the encrypted volume. The iManager Storage plug-in does not provide an encryption option. For information, see Creating an Encrypted Volume (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/bq2y6nb.html#bq2y6nb) in the *OES 2: Novell Storage Services File System Administration Guide* (http://www.novell.com/documentation/oes/nss_enu/index.html?page=/documentation/oes/nss_enu/data/front.html#front).

1a In iManager, expand the *Storage* role, then click *Volumes*.

1b If it is not already selected, select the archive server.

1c Click *New* to open the New Volume wizard.

1d Configure the new volume.

- Specify a name for the volume. For example, *ark*.
- Select the NSS pool you created in Section 5.3.3, "Creating an Archive Pool," on page 50, then select the *Allow the Volume to Grow to Pool Size* check box.
- Specify the desired attributes for the volume.

1e Click Finish.

2 (Optional, recommended) Create a directory in the archive volume where you want to store the archive database and archive data.

For detailed information, see Creating a Directory (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/bs3pfti.html) in the *OES 2: File Systems Management Guide* (http://www.novell.com/documentation/oes2/stor_filesys_lx_nw/index.html?page=/documentation/oes2/stor_filesys_lx_nw/data/hn0r5fzo.html).

2a Open your Web browser to the Novell Remote Manager Login page on the archive server, and then log in with your administrator username and password. For example, enter

```
https://192.168.1.1:8009
```

Replace *192.168.1.1* with the actual IP address or DNS name of your archive server.

2b Click *Manage Server* > *Volumes*.

2c Click the *Properties* icon next to the archive volume.

2d Type the name of the directory, then click *Create Subdirectory*.

3 Continue with Section 10.8, "Installing and Configuring the ArkSQL Server," on page 111.

10.7 Configuring the ArkSQL Configuration File

1 On the primary server in the cluster, configure the `sys:\arkManager\arkSQL.cnf` file.

For information, see Section 5.3.5, "Configuring the ArkSQL Configuration File," on page 52.

2 Copy the `sys:\arkManager\arkSQL.cnf` file from the primary server to the `sys:\arkManager` directory for each of the other servers in the cluster.

3 Continue with Section 10.8, "Installing and Configuring the ArkSQL Server," on page 111.

10.8 Installing and Configuring the ArkSQL Server

1. On each server in the cluster, install MySQL. Specify the shared volume location as the path of the archive database.

   For information, see Section 5.3.6, “Installing and Configuring the ArkSQL Server,” on page 53.
2. Continue with the next section, Configuring Archive Server Information.

10.9 Configuring Archive Server Information

1. On the primary server, configure the archive server information in the `sys:\arkManager\arkConfig.xml` file.

   For information, see Section 5.3.7, “Configuring Archive Server Information,” on page 56.
2. Copy the `sys:\arkManager\arkConfig.xml` file from the primary server to the `sys:\arkManager` directory for each of the others servers in the cluster.

   If you modify the settings for the archive server in the future, you must reconfigure the Basic elements in the `sys:\arkManager\arkConfig.xml` file and copy the modified file to all of the servers in the cluster.
3. Continue with Section 10.10, “Archiving File Versions,” on page 111.

10.10 Archiving File Versions

1. On the primary server in the cluster, configure the individual jobs for the archive cluster server. Specify the shared volume location as the path of the archive database.

   For information, see Chapter 7, “Configuring Jobs in ArkConfig,” on page 77.
2. Copy the `sys:\arkManager\arkConfig.xml` file from the primary server to the `sys:\arkManager` directory of the secondary servers in the cluster.
3. Continue with the next section, Editing the Autoexec.ncf File.

10.11 Editing the Autoexec.ncf File

When you install MySQL and ArkManager, commands are added to the `autoxec.ncf` file on the server to automatically start both programs when the NetWare server starts. Because Novell Cluster Services starts and stops the programs, these commands must be removed or commented out from each server in the cluster where they are installed.

1. On the primary server, open the `autoexec.ncf` file in a text editor.
2. Comment out or remove the following line:

   ```
   mysqld_safe --defaults-file=sys:\ArkManager\arkSQL.cnf
   ```

3. Comment out or remove the following line:

   ```
   arkstart
   ```

4 If you are using Port 3306 for your ArkSQL server, make sure the following line is commented out.

```
mysqld_safe --autoclose
```

This line should already be commented out as part of the install procedure.

5 If you are using another port for your ArkSQL server, make sure the following line is commented out:

```
mysqld_safe --port=3307 --autoclose
```

Replace *3307* with your actual port number. This line should already be commented out as part of the install procedure.

6 Save the `autoexec.ncf` file.

7 Copy the `autoexec.ncf` file to the other servers in the cluster.

8 Continue with Section 10.12, “Configuring Cluster Services to Automatically Stop and Start ArkManager and ArkSQL on System Reboot,” on page 112.

10.12 Configuring Cluster Services to Automatically Stop and Start ArkManager and ArkSQL on System Reboot

For each server in the cluster:

1 Set up Novell Cluster Services to start arkManager by default on reboot of the cluster.

1a Open ConsoleOne®.

1b In the ConsoleOne Properties dialog box, select *Scripts* > *Cluster Resource Load Script*.

1c Add the following commands to the end of the existing load script:

```
delay 2

mysqld_safe --defaults-file=sys:\ArkManager\arkSQL.cnf

search add sys:\arkManager

arkstart
```

1d Click *Apply*.

2 Set up Novell Cluster Services to stop arkManager services by default on the server down command.

2a In the ConsoleOne Properties dialog box, click *Scripts* > *Cluster Resource Unload Script.*

2b Add the following command to the beginning of the existing unload script:

```
arkstop
```

2c Click *Apply*, then click *Close*.

3 Restart each of the servers in the cluster to offline and then online the resources.

For information, see the *NW65 SP8: Novell Cluster Services 1.8.5 Administration Guide (http://www.novell.com/documentation/oes2/clus_admin_nw/index.html?page=/documentation/oes2/clus_admin_nw/data/h4hgu4hs.html#h4hgu4hs.html)*.

When the servers come back up, Novell Archive and Version Services is running on the primary server. The secondary servers will be on hot standby, waiting to be called to action.

Coexistence and Migration Issues for Archive and Version Services

11

This section discusses the issues involved in the coexistence and migration of Novell® Archive and Version Services for NetWare®.

11.1 Compatibility with Operating Systems and Services

The following table summarizes the compatibility of Archive and Version Services with various network operating systems and directory services versions, along with any dependencies on other products or services.

Table 11-1 *Compatibility of the Archive and Version Services and Other OES Components*

NetWare Operating System	OES SP1 NetWare and NetWare 6.5 SP4
Linux Operating System	OES 2 Linux
Directory Services	Novell eDirectory™ 8.7.3 or later
Management Services	Novell iManager 2.7
Dependencies	◆ Novell Storage Services™ file system on OES SP1 NetWare or later. Target source volumes can be NetWare 6.5 or later NSS volumes. ◆ NCP™ or Native File Access Protocols (AFP, CIFS, NFS) for user authentication and access to file versions in the archive database ◆ MySQL ◆ Novell Archive and Version Services File Version Utility for Windows 2000/XP ◆ NetStorage (Archive function) for browser-based access to file versions

11.2 File System Support for Source Data

The source and archive data volumes must be Novell Storage Services data volumes that use the OES NetWare and NetWare 6.5 media format. For information, see Section 4.5, “Storage Media Prerequisites and Guidelines,” on page 39.

An archive server and its source servers where NSS data volumes reside can run either operating system in a mixed environment. Make sure the servers are running the latest upgrades.

The following table summarizes the compatibility of Novell Archive and Version Services 2.1 with various network operating systems versions and file systems.

Table 11-2 *Support for Potential Target Source Volumes*

Target File System	Target Operating System	Supported (Yes/No)
NSS volumes	OES 2 Linux	Yes
NSS volumes	OES NetWare	Yes
NSS volumes	NetWare 6.5	Yes
NSS volumes	NetWare 6.0 and earlier	No
NSS volumes	OES Linux	No
NetWare Traditional volumes	All NetWare versions	No
NCP volumes (traditional Linux volumes with NCP Server)	OES Linux	No
Traditional Linux volumes	OES Linux or SLES	No

11.3 UTF-8 Encoding for Filenames

ArkManager uses UTF-8 encoding for filenames. Either the Language Code pages of the NCP client and server must be the same, or you must use a current version of the Novell Client to provide UTF-8 support and enable the UTF-8 service in the client. CIFS clients have included UTF-8 encoding support since NetWare 6.0 SP 2 and NetWare 6.5. For information, see Section 9.11, “Enabling UTF-8 Encoding Support for Clients,” on page 104.

11.4 Compatibility Issues for the Archive Versioning iManager Plug-In

11.4.1 Novell iManager 2.7

The Archive Versioning plug-in for OES 2 NetWare requires Novell iManager 2.7. You cannot use this version of the plug-in with earlier releases of iManager. For information, see Upgrading to iManager 2.5 in the Novell iManager 2.7 Administration Guide (http://www.novell.com/documentation/imanager27/imanager_admin_27/index.html?page=/documentation/imanager27/imanager_admin_27/data/hk42s9ot.html#hk42s9ot) .

11.4.2 Interoperability of the Archive Versioning Plug-In

The Storage, Archive Versioning, File Protocols (NetWare), and Cluster Services plug-ins share common code. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins. If you remove only one of the plug-ins, it removes the common code and breaks the remaining installed plug-ins.

11.4.3 Web Browser Language Setting

The Archive Versioning iManager plug-in might not operate properly if the highest priority (top position) Language setting for your Web browser is set to a language other than one of the supported languages. To avoid problems, in your Web browser, click *Tools* > *Options* > *Languages*, and then set the first language preference in the list to a supported language.

11.4.4 Protocols

The following table provides information about the protocols needed to use iManager to manage storage in a heterogeneous environment. A protocol annotated with an asterisk is the default and is automatically configured on the servers. The protocols you use must be loaded and running on both the iManager server and the target server you want to manage.

Table 11-3 *Protocols Needed to Use the Archive Versioning Plug-In to Manage Archive Servers*

iManager 2.5 Server	**Target Archive and Version Services Server**		
Operating System	**OES NetWare and Later, or NetWare 6.5 SP4**	**NetWare 6.5 SP3**	**NetWare 6.5 SP2**
OES Linux and Later	WBEM*	WBEM (Start OpenWBEM.)	
	CIFS	CIFS	CIFS (Field Patch 2B)
OES NetWare and Later, or NetWare 6.5 SP4	WBEM*	WBEM (Start OpenWBEM.)	
	NCP	NCP*	NCP*
	CIFS	CIFS	CIFS (Field Patch 2B)
NetWare 6.5 SP3	WBEM*	WBEM (Start OpenWBEM.)	
	NCP	NCP*	NCP*
	CIFS	CIFS	CIFS (Field Patch 2B)
NetWare 6.5 SP2	NCP*	NCP*	NCP*

* Marks the default protocol that is automatically configured.

() Additional requirements appear in parens.

WBEM

Where WBEM is the default protocol, OpenWBEM is loaded and runs automatically when you start the server. Otherwise, you must start OpenWBEM to use the protocol.

1 At the server console, enter

```
openwbem
```

If you receive file protocol errors, it might be because OpenWBEM is not running. For information about installing OpenWBEM, see OES 2: OpenWBEM Services Administration Guide (http://www.novell.com/documentation/oes2/mgmt_openwbem_lx_nw/index.html?page=/documentation/oes2/mgmt_openwbem_lx_nw/data/front.html#front).

CIFS

Where it is available, CIFS must be configured before you can use it. An additional CIFS setup requirement for Field Patch 2B is noted where it is required. For information, see the OES 2: Native File Access Protocols Guide (http://www.novell.com/documentation/oes2/file_afp_cifs_nfs_nw/index.html?page=/documentation/oes2/file_afp_cifs_nfs_nw/data/h9izvdye.html#h9izvdye).

NCP

NCP is the default protocol when the iManager server and target server are both NetWare 6.5 SP3 or NetWare 6.5 SP2.

Security Considerations for Archive and Version Services

12

This section discusses the following security issues for Novell® Archive and Version Services for NetWare®:

12.1 NSS Encrypted Volumes

If you plan to version files from one or more encrypted source volumes, the archive volume must also be an encrypted volume. Otherwise, the data is stored nonencrypted on the archive volume.

For information about NSS Encrypted Volume Support, see Managing Encrypted NSS Volumes in the OES 2: NSS File System Administration Guide (http://www.novell.com/documentation/oes2/stor_nss_lx_nw/index.html?page=/documentation/oes2/stor_nss_lx_nw/data/front.html#front).

The data is not secure when transported between the source volume and the archive database. For information, see Section 12.2, “Secure Data Transfer,” on page 119.

12.2 Secure Data Transfer

Because the archive server moves data in nonencrypted format, the archive server should be located behind the corporate firewall. If you are versioning data on remote servers, use a virtual private network (VPN) connection between the two.

12.3 eDirectory and MySQL Administrator Passwords

The passwords for the Novell eDirectory™ ArkManager administrator user and the MySQL database administrator user are passed to ArkManager in the `sys:\arkManager\arkConfig.xml` file the first time ArkManager runs as part of the basic server information. For information, see Section 5.3.7, “Configuring Archive Server Information,” on page 56.

Thereafter, if you modify the password for the ArkManager administrator user in eDirectory or the password for the database administrator user in MySQL, you must pass the new password to ArkManager by adding them again to the `arkConfig.xml` file. For information, see Section 9.7, “Updating Passwords in ArkManager,” on page 103.

12.4 MySQL Secure Installation

You are strongly encouraged to use a secure installation of MySQL for production environments. With the default installation instead of a secure installation, any local user can connect without a password and be treated as the anonymous user. This creates a high security risk in your production environment.

12.5 User Authentication

Archive and Version Services uses eDirectory to authenticate users who access the archive database. For more information, see Section 4.1, “Network Architecture Prerequisites and Guidelines,” on page 37.

A user assigned with read and write permissions on NetWare Archive volume can modify the versions of the files, using NSS File Version Utility.

12.6 Archive Database and Data Backup

You should store the archive database and archive data in a directory in the archive volume for easy backup. The archive data is always stored in subdirectories at the base of the archive path. For information, see the following:

- Section 9.8, “Backing Up the Archive Database,” on page 103
- Section 9.9, “Backing Up the Archive Data,” on page 103
- Section 9.10, “Recovering the Archive Database,” on page 104

XML Elements and Attributes for ArkConfig

A

This section defines the XML elements and attributes that you use to configure versioning jobs for your Novell® Archive and Version Services 2.1 for NetWare® server. The XML elements have a particular hierarchy that you must keep in mind as you define the Basic element, Defaults element, and multiple Job elements. See Figure A-1 on page 122 to understand the parent-child relationships between XML elements defined for the `sys:\arkManager\arkConfig.xml` file.

The following table defines the XML elements and attributes that you use in the `sys:\arkManager\arkConfig.xml` file. Some elements appear first as children of elements in the next-higher level of the hierarchy and as parents to their own set of child elements in the next-lower level of the hierarchy.

The examples provided for the elements are sample code to illustrate how the configured properties appear in the XML file. You must modify the properties to meet your specific needs as you create your `arkConfig.xml` file.

At a minimum, you must configure the authentication information in `arkConfig.xml` for ArkManager to run. For a sample of the basic configuration, see Section B.2, “Sample of a Basic Configuration for ArkConfig.xml,” on page 144.

For a sample of the full configuration, see Section B.3, “Sample of a Full Configuration for ArkConfig.xml,” on page 146.

Figure A-1 *Hierarchical Parent-Child Relationships between XML Elements Used in the arkConfig.xml File*

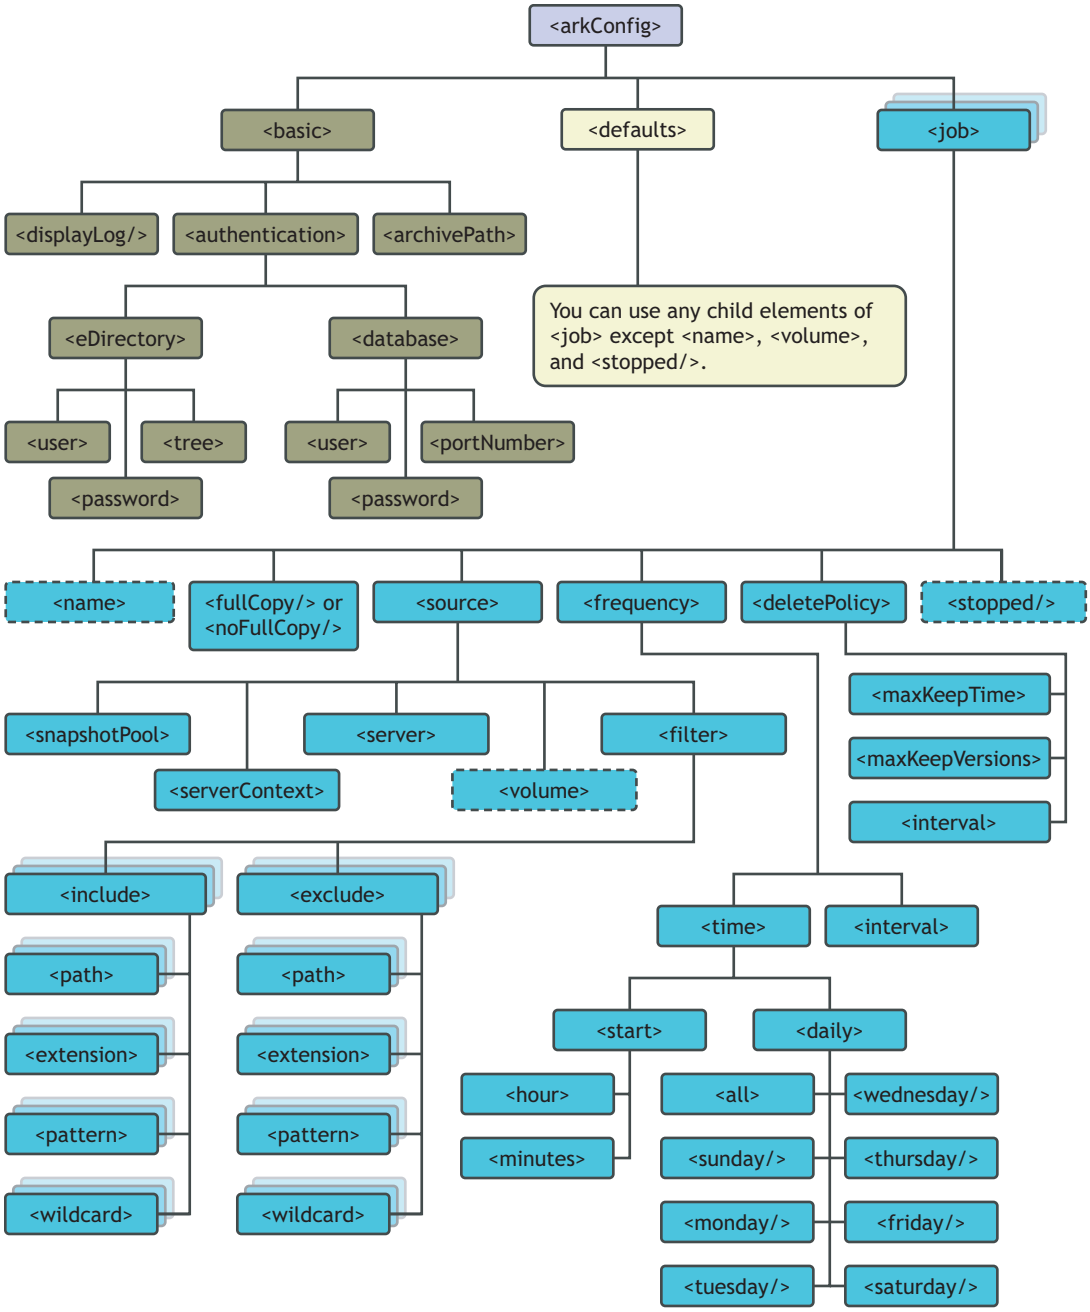

Table A-1 *Elements of the arkConfig.xml File*

Parent Element	Child Elements and Attributes	Description
<arkConfig> (Root element)	<basic>	A single Basic tag set surrounds the management information that applies to all jobs. Its child elements specify where to display the log, the authentication information, and the storage location where the archive data (versions) reside. At a minimum, configure the <authentication> information for ArkManager 2.1 to run. For information about its child elements, see <basic> in the Parent Element column of this table.
	<defaults>	A single Defaults tag set surrounds the child elements that define the default settings to use if the elements are not set in individual jobs. Each job uses the values configured in the Defaults element only for those values that are not otherwise set for that job. For information about its child elements, see <defaults> in the Parent Element column of this table.
	<job>	A Job tag set surrounds the child elements that define the settings for a given job. Create one job for each unique information set that you want to version. An archive server can version a given source volume with only one job. If you want to create multiple jobs for a source volume, you must create the jobs on different archive servers. The values configured in a given job override the default settings only for that job. If a value is not set in the Job element, the default settings apply for only that element. Use the attributes optionally to indicate whether the job is to be added, modified, or deleted from the archive server. For information about its child elements, see <job> in the Parent Element column of this table.

Parent Element	Child Elements and Attributes	Description
<basic> (Child to <arkConfig>)	<displayLog/>	(Optional) If this tag is included in the arkConfig.xml file, it prints log entries for all jobs to the server's console screen in addition to their normal display in the Archive / Version Log. This is an empty element. It has no child elements and defines no data values. Either the shortened version or the long version of this element is valid. For example, either of the following formats work: <displayLog/> <displayLog></displayLog>
	<authentication>	Contains child elements that provide information about the administrator user. This user must have rights to all servers being accessed by this archive server. Child elements also define the ArkSQL database user, password, and port number. The ArkSQL user needs server-specific ArkSQL rights. The <authentication> information must be set before ArkManager can run. For information about its child elements, see <authentication> in the Parent Element column of this table.
	<archivePath>	Specifies the local path to the NSS volume in the archive server where the archive database of file versions is stored. <archivePath> *volume_name:\directory* </archivePath> Placing a backslash between the colon and the directory name, as shown in in this example, is optional. If you are setting up an active-passive clustered solution for the archive server using Novell Cluster Services™, make sure to specify the path to the virtual storage location.

Parent Element	Child Elements and Attributes	Description
<defaults> (Child to <arkConfig>)	<fullCopy/> <noFullCopy/> <source> <frequency> <deletePolicy>	For information about these child elements, see the Job element (<job>) in this table. The Defaults element can contain any tags that you use in a Job element, except for <name>, <volume>, and <stopped/>. Although you can have multiple Job elements, only a single Defaults element appears in the arkConfig.xml file.

Parent Element	Child Elements and Attributes	Description
<job> (Child to <arkConfig>)	<name>	Specifies the unique name of this versioning job. This name can contain spaces. You might want the job name to reflect information about the data that is being versioned. For example, if you have a job that versions data for a user1 volume on the svr1 server, you might name the job: <name>*user1[svr1]*</name> Use the <name> tag as a child only of the Job element; it is invalid in the Defaults element.
	<fullCopy/>	(Optional) If used, sets a flag so that the job copies all files that pass the filter from the source volume to the archive database the first time the job runs or if a full copy runs but did not finish previously. In any given Job element or Defaults element, use <fullCopy/> or use <noFullCopy/>, not both. If neither one is specified, <fullCopy/> is the default software setting. This is an empty element; it has no child elements and no data values.
	<noFullCopy/>	(Optional) If used, sets a flag so that the job does not copy all files that pass the filter from the source volume to the archive database the first time the job runs. This is an empty element; it has no child elements and no data values.
	<source>	Contains child elements that define the specific server context, server, paths, directories, or file types to be archived for this job. For information about its child elements, see <source> in the Parent Element column of this table.
	<frequency>	Contains child elements that define the start time and intervals for versioning in this job. For information about its child elements, see <frequency> in the Parent Element column of this table.

Parent Element	Child Elements and Attributes	Description
<job> (cont'd)	<deletePolicy>	Contains child elements that specify when to delete a version, such as by age of the version (elapsed time of creation), by the number of versions that exist (oldest deleted first), or by the interval of time before the next versioning process. For information about its child elements, see <deletePolicy> in the Parent Element column of this table.
	<stopped/>	Specify this property to define the job but leave it in a Stopped state until you manually activate the job. If the Stopped parameter is not used, the job starts, according to the Run Schedule settings, the next time ArkManager runs. Use the <stopped/> tag as a child only of the Job element; it is invalid in the Defaults element.
<authentication> (Child to <basic>)	<eDirectory>	Contains child elements that specify the Novell eDirectory™ authentication information, including the administrator user, the password, and the eDirectory tree where the administrator user, archive server, and source servers reside. For information about its child elements, see <eDirectory> in the Parent Element column of this table.
	<database>	Contains child elements that contain information about the archive server's ArkSQL database. This user must have server-specific rights to the ArkSQL server. For information about its child elements, see <database> in the Parent Element column of this table.

Parent Element	Child Elements and Attributes	Description
<eDirectory> (Child to <authentication>)	<user>	Specifies the Novell eDirectory Common Name of the administrator user who has the appropriate rights to the original data location and to the archive server's data location. For example: <user>*admin.server_context*</user>
	<password>	Specifies the password for the specified administrator user. For example: <password>*admin_pwd*</password> The first time that you run arkManager, the arkConfig.xml file must contain a <password> element with the correct password for the administrator user. If arkManager finds a <password> element, it reads the password, encrypts it, then stores the encrypted password in a local store. It removes the password from the arkConfig.xml file, then saves the file. **IMPORTANT:** The password is stored initially in arkConfig.xml unencrypted. Until the password is moved to an encrypted store, the arkConfig.xml file is protected from unauthorized access by the NSS file system rights that apply to the sys:\arkManager directory and the arkConfig.xml file itself. The <password> element is not necessary unless you need to change the saved password. ArkManager refers to the encrypted password whenever you use the arkStart command. If you change the administrator user's password on the network, you must also change the password that arkManager uses when it starts. ArkManager cannot restart until you pass it the correct password in the arkConfig.xml file. If you need to change the password, stop ArkManager (arkstop), add an eDirectory <password> element to the arkConfig.xml file with the new password, then start (arkstart) the ArkManager.
	<tree>	Specifies the tree where the administrator user, archive server, and source servers exist. For example: <tree>*site1_tree_name*</tree> <tree>eur_tree</tree>

Parent Element	Child Elements and Attributes	Description
<database> (Child of <authentication>)	<user>	The username of the ArkSQL database administrator. For example: <user>*root*</user>
	<password>	The password of the ArkSQL database administrator. For example: <password>*arksql_password*</password> The first time that you run arkManager, the arkConfig.xml file must contain a <password> element with the correct password for the MySQL root user. If arkManager finds a <password> element, it reads the password, encrypts it, then stores the encrypted password in a local store. It removes the password from the arkConfig.xml file, then saves the file. **IMPORTANT:** The password is stored initially in arkConfig.xml unencrypted. Until the password is moved to an encrypted store, the arkConfig.xml file is protected from unauthorized access by the NSS file system rights that apply to the sys:\arkManager directory and the arkConfig.xml file itself. The <password> element is not necessary unless you need to change the saved password. ArkManager refers to the encrypted password whenever you use the arkStart command, which also starts ArkSQL. If you change the root user's password for MySQL, you must also change the database password that arkManager uses when it starts. ArkManager cannot restart until you pass it the updated password in the arkConfig.xml file. If you need to change the password, stop ArkManager (arkstop) and ArkSQL , add a database <password> element to the arkConfig.xml file with the new password, then start (arkstart) the ArkManager.
	<portNumber>	The port number used for ArkSQL database communications. For example: <portNumber>*3306*</portNumber>

Parent Element	Child Elements and Attributes	Description
<source> (Child to <defaults> and to <job>)	<snapshotPool>	Specifies the name of the destination pool where a snapshot of the source volume is created temporarily at the beginning of each versioning. For example: <snapshotPool>*pool_name* </snapshotPool> Snapshots make it possible to save point-in-time versions of all eligible source files, even if a file is open at the time. Eligible files are those files that exist at the end of the epoch, changed during the epoch, and met the filter requirements. By default, if no snapshot pool is specified, ArkManager copies the eligible source files directly from the source volume. Exclusively open files cannot be versioned and data might be inconsistent. If the specified snapshot pool is inactive or otherwise not available when a versioning job begins, the job copies the files directly from the source volume.
	<serverContext>	Specifies the unique name of the context of the server that hosts the source volume, which contains the data that the job versions and saves to the archive database. Type the server context in the Novell common dot format, from lowest to highest level. It does not include the tree. For example: <serverContext> grp1.dept1.site1.examplecompany </serverContext> <serverContext> sales.mktg.eur.acme </serverContext>
	<server>	Specifies the unique name of the server in the specified context where the original data is stored. For example, to specify a server named srv1: <server>*srv1*</server> Use the <server> tag as a child only of the Job element; it is invalid in the Defaults element.

Parent Element	Child Elements and Attributes	Description
<source>	<volume>	Specifies the unique name of the NetWare 6.5 or later NSS volume in the specified source server where the original data is stored.
(cont'd)		For example, to specify user1 as the source volume: <volume>*user1*</volume> Do not place a colon after the name of the source volume. Use the <volume> tag as a child only of the Job element; it is invalid in the Defaults element.
	<filter>	(Optional) Contains child elements and attributes that specify the criteria for filtering data so that only selected data is versioned.
	<filter overrideDefaults="false">	Use the attributes optionally to indicate whether the filter is to be used in combination with filters set in the Defaults element (false) or if the filter is to be used in place of the filters set in the Defaults element (true).
	<filter overrideDefaults="true">	For information about its child elements, see <filter> in the Parent Element column of this table.
<filter>	<include>	(Optional) Contains child elements that specify traits of data you want to include in the job.
(Child to <source>)		For information about its child elements, see <include> in the Parent Element column of this table.
	<exclude>	(Optional) Contains child elements that specify traits of data you want to exclude in the job.
		For information about its child elements, see <exclude> in the Parent Element column of this table.

Parent Element	Child Elements and Attributes	Description
<include> (Child to <filter>)	<path>	(Optional) Specifies the relative path of directories in the specified volume that you want to include in the versioning process. If you want to include eligible files only in the specified paths, make sure to exclude the root path in an exclude path statement. For example, to specify a relative path of the volume_name:\dept1\data directory that contains data you want to include: <include> <path>*dept1**data*\</path> </include> Repeat this element for each directory path in the specified volume that contains data that you want to version.
	<extension>	(Optional) Specifies the extension of files in the specified volume that you want to include in the versioning process. Use the preceding dot followed by the characters of the file extension. Repeat this element for each file extension type that you want to version. For example, to specify that you want to version only files with no extension (.) and files with .xxx and .yyy extensions: <include> <extension>.</extension> <!-- No extension --> <extension>.*xxx*</extension> <extension>.*yyy*</extension> </include>
	<pattern>	(Optional) Specifies the regular expression pattern to match for files that you want to include in the versioning process. For example, to set criteria to include only files with names that start with the letter "a". <include> <pattern>.**a*.*</pattern> </include> Repeat this element for each pattern that you want to match to identify files for versioning. This element does not support the full set of PERL regular expressions. For more information, see "Pattern Elements" on page 34.

Parent Element	Child Elements and Attributes	Description
<include> (cont'd)	<wildcard>	(Optional) Functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to include files that start with d of type .sxi: <include> <wildcard>d*sxi</wildcard> </include>
<exclude> (Child to <filter>)	<path>	(Optional) Specifies the relative path of directories in the specified volume that you want to exclude from the versioning process. For example, to specify a relative path of the volume_name:\dept1\apps directory that contains application data you want to exclude from versioning: <exclude> <path>*dept1**apps*\</path> </exclude> Repeat this element for each directory path in the specified volume that contains data that you do not want to version.
	<extension>	(Optional) Specifies the extension of files in the specified volume that you want to exclude from the versioning process. Use the preceding dot followed by the characters of the file extension. Repeat this element for each file extension type that you do not want to version. For example, to specify that you want to exclude .zzz files from versioning: <exclude> <extension>.*zzz*</extension> </exclude>

Parent Element	Child Elements and Attributes	Description

<exclude>	<pattern>	(Optional) Specifies the regular expression pattern to match for files that you want to exclude from the versioning process.
		
		For example, to set criteria to exclude files with names that start with the letter "a":
		
		`<exclude>`
		`  <pattern>.*\\a.*</pattern>`
		`</exclude>`
		
		Repeat this element for each pattern that you want to match to identify files to be excluded from versioning.
		
		This element does not support the full set of PERL regular expressions. For more information, see "Pattern Elements" on page 34.

(cont'd)	<wildcard>	(Optional) Functions like a wildcard in directory searches. Replace characters with an asterisk (*) to search for files that match. For example, to exclude files that start with d of type .tmp:
		
		`<exclude>`
		`  <wildcard>d*tmp</wildcard>`
		`</exclude>`

Parent Element	Child Elements and Attributes	Description

<frequency>
 	
(Child to <defaults> and to <job>)
	<time>	(Conditional) Contains child elements that specify the start time, based on a 24-hour local time scale, of jobs that are scheduled to occur on specified days of the week. A valid hour entry is an integer value ranging from 0 to 23.
		
		Use either the Time element or Interval element, but not both in the same Job element.
		
		For example, to start the versioning at 11:15 p.m. (23:45) every day of the week:
		
		`<frequency>`
		
		`  <time>`
		
		`    <start>`
		`      <hour>23</hour>`
		`      <minute>15</minute>`
		`    </start>`
		
		`    <daily>`
		`      <all/>`
		`    </daily>`
		
		`  </time>`
		
		`</frequency>`
		
		For information about its child elements, see <time> in the Parent Element column of this table.

	<interval>	(Conditional) Specifies the elapsed time between scheduled versions in seconds, minutes, hours, or days.
		
		You must use one of the attributes in the Interval tag to specify the units of the integer value you specify as the interval of time to wait between when the versioning begins. For example, to specify an interval of 30 minutes:
		
		`<interval unit="minutes">30`
		`</interval>`
		
		You must specify a non-zero value for the interval or the job cannot start. If no interval value is set, 0 (zero) is the default setting.
		
		If no interval unit is set, seconds are the default unit setting.
		
		Use either the Time element or Interval element,

	<interval unit="seconds">	
	<interval unit="minutes">	
	<interval unit="hours">	
	<interval unit="days">	

Parent Element	Child Elements and Attributes	Description
<time> (Child to <frequency>)	<start>	Contains child elements that specify the hour and minutes that the job is scheduled to run on the specified days of the week. For information about its child elements, see <start> in the Parent Element column of this table.
	<daily>	Contains child elements that specify the days of the week that the job is scheduled to run. You must specify at least one day or the job cannot start. For information about its child elements, see <daily> in the Parent Element column of this table.
<start> (Child to <time>)	<hour>	Specifies the hour of the day that the job begins, based on a 24-hour clock. A valid hour entry is an integer value ranging from 0 to 23. For example, to set the hour to 10 p.m.: <hour>22</hour> If no hour is set, midnight (00) is the default hour setting.
	<minute>	Specifies the minutes of the hour in the day that the job is scheduled to begin. A valid minutes entry is an integer value ranging from 0 to 59. For example, to set the minutes after the hour to 45 minutes: <minute>45</minute> If no minute value is set, zero (00) is the default minute setting.

Parent Element	Child Elements and Attributes	Description
<daily> (Child to <time>)	<monday/> <tuesday/> <wednesday/> <thursday/> <friday/> <saturday/> <sunday/> <all/>	Specify one or more of the child elements as the days of the week you want to schedule the versioning to occur. If no <daily> value is set, the job cannot start. For example, to specify versioning to occur on Tuesday, Friday, and Sunday of each week: <daily> <tuesday/> <friday/> <sunday/> </daily> Use the <all/> tag to specify all seven days of the week. For example: <daily> <all/> </daily> Each of the <daily> child elements is an empty element; it has no child elements and no data values.

Parent Element	Child Elements and Attributes	Description

<deletePolicy>
 	
(Child to <defaults> and to <job>)
	<maxKeepTime>
	
	<maxKeepTime unit="seconds">
	
	<maxKeepTime unit="minutes">
	
	<maxKeepTime unit="hours">
	
	<maxKeepTime unit="days">
	
	<maxKeepTime keepCurrentCopy="true">
	
	<maxKeepTime keepCurrentCopy="false">
		Use one of the unit attribute values in the <maxKeepTime> tag to specify the units of the integer value you specify as maximum lifetime of a version.
		
		If no maxKeepTime value is set, file versions are retained indefinitely.
		
		If a value is specified without a unit attribute, “seconds” is the default unit of the value specified.
		
		Use the keepCurrentCopy attribute to designate that at least one file version of current files remains in the database, even if the maxKeepTime elapses.
		
		If keepCurrentCopy is set to True (default), the archive keeps an existing file version as long as its source file is current on the source volume, beyond the maxKeepTime. After the user deletes the current source file, the deletion is noted at the next scheduled epoch. If the file version’s age is within the maxKeepTime, the archive database retains a copy of the file version until its maxKeepTime elapses. When the file version’s age exceeds the maxKeepTime, the archive deletes the file version at the next scheduled delete time.
		
		If keepCurrentCopy is set to False, the archive deletes a file version when it exceeds the maxKeepTime, even if the file version is the only one in the archive database for a given source file, and whether the source file is current or deleted.
		
		The keepCurrentCopy attribute is optional. If the keepCurrentCopy attribute is not otherwise specified, the default value is True.
		
		For example, to keep versions for 90 days before purging, to keep at least the most current version, and to schedule the purging of versions with 1-hour intervals:
		
		`<deletePolicy>`
		
		`<maxKeepTime unit="days" keepCurrentCopy="true">90 </maxKeepTime>`
		
		`<interval unit="hours">1 </interval>`
		
		`</deletePolicy>`

Parent Element	Child Elements and Attributes	Description
<deletePolicy> (cont'd)	<maxKeepVersions>	Specifies the maximum number of versions of each file to keep in the archive. As the number of versions exceed this integer value, the oldest version is deleted. For example, to keep up to 1000 versions of each file before purging and to schedule the purging of versions to begin at the default interval of 24 hours before the next scheduled versioning process: <deletePolicy> <maxKeepVersions>1000 </maxKeepVersions> </deletePolicy>
	<interval> <interval unit="seconds"> <interval unit="minutes"> <interval unit="hours"> <interval unit="days">	The interval represents the amount of time to wait from the time a Delete process ends until another Delete process begins. If a value is not specified, 24 hours is the default delete policy interval. The time involved in deleting file versions varies with the amount and complexity of data stored in the archive server. The Delete Schedule operates separately from the Version Schedule. For example, suppose you set the Delete Schedule to 2 days. When you activate the job, the Delete process begins. When it is done, it sets an interval timer. After two days elapse, the Delete process runs. The timer is inactive while the process runs. When the delete process ends, the interval timer begins again. The process repeats until the job is deactivated. For example, to keep up to 100 versions of each file before purging and to schedule the purging of versions with a 2-hour interval: <deletePolicy> <maxKeepVersions>100 </maxKeepVersions> <interval unit="hours">2 </interval> </deletePolicy> If an attribute is not specified for the <interval> tag, "seconds" is the default unit of the value specified.

Sample Configuration Files

B

This section illustrates the use of XML and CNF files used to configure Novell® Archive and Version Services 2.1 for NetWare®. For your convenience, all sample files can be found in the `sys:\arkManager` directory of your archive server. Use the following table to determine which samples to use.

Table B-1 *Possible Tasks*

If You Need To	View This Sample File	Sample File (sys:\arkManager\)
Configure `ArkSQL.cnf` with the MySQL database properties to work with ArkManager 2.1.	Sample of the Database Configuration for ArkSQL.cnf (page 141)	`arkSQL_sample.cnf`
Configure authentication information for the archive server. This is the minimum configuration information in the `arkConfig.xml` file that is needed to use the Archive Versioning plug-in for Novell iManager. Of course, you must continue to define the other elements to run jobs.	Sample of a Basic Configuration for ArkConfig.xml (page 144)	`arkConfig_sample_basic.xml`
Configure all properties for the archive server, individual jobs, and default job settings in the `arkConfig.xml` file.	Sample of a Full Configuration for ArkConfig.xml (page 146)	`arkConfig_sample_full.xml`

B.1 Sample of the Database Configuration for ArkSQL.cnf

Novell Archive and Version Services 2.1 uses the `ArkSQL.cnf` file to configure the ArkSQL settings. Use the `sys:\arkManager\arkSQL_sample.cnf` file as a guide to configure the `sys:\arkManager\arkSQL.cnf` file.

For information about configuring `ArkSQL.cnf`, see Section 5.3.6, “Installing and Configuring the ArkSQL Server,” on page 53.

```
#This is the sample file to configure the MySQL server instance for ArkManager
use.


#Copy the contents to sys:\arkManager\arkSQL.cnf, and modify it accordingly.
```

```
# The MySQL server

[mysqld]

port        = 3306

datadir         = Ark:/archive

default-table-type=InnoDB

lower_case_table_names=1

skip-locking

set-variable    = key_buffer=32M

set-variable    = max_allowed_packet=5M

set-variable    = table_cache=512

set-variable    = sort_buffer=2M

set-variable    = read_buffer_size=1M

set-variable    = net_buffer_length=16K

set-variable    = query_cache_size=32M

#log-bin
```

```
server-id    = 1

thread_cache_size = 8

# thread cache

# thread concurrency

# Point the following paths to different dedicated disks

#tmpdir         = /tmp/

#log-update     = /path-to-dedicated-directory/hostname  # deprecated

# Uncomment the following if you are using BDB tables

#set-variable   = bdb_cache_size=4M

#set-variable   = bdb_max_lock=10000

# Comment the following if you are using Innobase tables

#skip-innodb

# Uncomment the following if you are using Innobase tables

innodb_data_file_path = ibdata1:400M:autoextend

#innodb_data_home_dir = Ark:/archive/
```

```
innodb_log_group_home_dir = Ark:/archive/

innodb_log_arch_dir = Ark:/archive/

set-variable = innodb_mirrored_log_groups=1

set-variable = innodb_log_files_in_group=3

set-variable = innodb_log_file_size=5M

set-variable = innodb_log_buffer_size=8M

innodb_flush_log_at_trx_commit=1

innodb_log_archive=0

set-variable = innodb_buffer_pool_size=16M

set-variable = innodb_additional_mem_pool_size=2M

set-variable = innodb_file_io_threads=4

set-variable = innodb_lock_wait_timeout=50

set-variable = transaction_isolation=READ_COMMITTED
```

B.2 Sample of a Basic Configuration for ArkConfig.xml

Before ArkManager 2.1 can run, you must configure the `sys:\arkManager\arkConfig.xml` file with basic authentication information for Novell eDirectory™ and MySQL. Use the `sys:\arkManager\arkConfig_sample_basic.xml` file as a guide for configuring the minumum set of elements needed to run ArkManager.

Use the full example in the `sys:\arkManager\arkConfig_sample_full.xml` file as a guide to configure default jobs settings and one or more individual versioning job. For information, see Section B.3, “Sample of a Full Configuration for ArkConfig.xml,” on page 146

For information about the elements and their usage, see Chapter A, “XML Elements and Attributes for ArkConfig,” on page 121.

```
<!-- This is the sample file to configure the Basic elements for
     arkConfig.xml. Basic element configure authentication for
     the archive server and the MySQL server instance. Copy it
     to sys:\arkManager\arkConfig.xml, and modify it accordingly.


-->


<arkConfig> <!-- Root element -->


 <basic> <!-- Start Basic element -->


   <displayLog/>


   <authentication> <!-- Administrator and database authentication
                        information -->


     <eDirectory> <!-- eDirectory Administrator User information -->


       <user>admin.xyz_inc</user> <!-- Use the eDirectory Common Name.
                                  -->


       <password>pass_word</password>


       <tree>archive_tree</tree>


     </eDirectory>


     <database> <!-- MySQL Administrator User information -->
```

```
        <user>root</user>

        <password>pass_word</password>

        <portNumber>3306</portNumber>

      </database>

   </authentication>

   <archivePath>archive_volume:</archivePath>

 </basic> <!-- End Basic element -->

</arkConfig>
```

B.3 Sample of a Full Configuration for ArkConfig.xml

The `sys:\arkManager\arkConfig.xml` file contains basic settings, default job settings, and individual job settings. Use the `sys:\arkManager\arkConfig_sample_full.xml` file as a guide for configuring jobs for your archive server.

For information about the elements and their usage, see Chapter A, “XML Elements and Attributes for ArkConfig,” on page 121.

```
<!-- This is the sample file to configure jobs manually. Copy it
     to sys:\arkManager\arkConfig.xml, and modify it accordingly.
-->


<arkConfig> <!-- Root element -->


 <basic> <!-- Start Basic element -->


   <displayLog/>
```

```
  <authentication> <!-- Administrator and database authentication
                       information -->

    <eDirectory> <!-- eDirectory Administrator User information -->

      <user>admin.xyz_inc</user> <!-- Use the eDirectory Common Name.
                                  -->

      <password>pass_word</password>

      <tree>archive_tree</tree>

    </eDirectory>

    <database> <!-- ArkSQL Administrator User information -->

      <user>root</user>

      <password>pass_word</password>

      <portNumber>3306</portNumber>

    </database>

  </authentication>

  <!-- Location of the archive of versioned files -->

  <archivePath>archive_volume:</archivePath>

</basic> <!-- End Basic element -->
```

```
<defaults> <!-- Start Defaults element -->

  <source> <!-- Location of the data that you want to version -->

    <serverContext>xyz_inc</serverContext>

    <filter>

    <!-- By default, everything on the source volume is eligible for
         file versioning. This example filter includes everything on
          the volume (by default), except for the files with extension
         .tmp, or under \dir1\dir2 directory, or matches pattern
         .*[uU]2.* (such as 1U2.exe), or matches wildcard abc*
         (such as abcd.c).
    -->

      <exclude>

        <extension>.tmp</extension>

        <path>\dir1\dir2</path>

        <pattern>.*[uU]2.*</pattern>

        <wildcard>abc*</wildcard>

      </exclude>

    </filter>

  </source>

  <frequency> <!-- Schedule for versioning data -->
```

```
    <time>
      <start>
        <hour>12</hour>
        <minute>30</minute>
      </start>
      <daily>
        <saturday/>
      </daily>
    </time>
  </frequency>
</defaults> <!-- End Defaults element -->
<job> <!-- Begin a Job element -->
  <name>server1_volume1</name>
  <frequency>
    <interval unit="seconds">20</interval>
```

```
</frequency>

<source>

  <server>server1</server>

  <volume>volume1</volume>

  <snapshotPool>snapPool</snapshotPool> <!-- (Optional) Use NSS pool
                                            snapshots. -->

  <filter overrideDefault="true"> <!-- Override default filter. -->

  <!-- This filter excludes everything on the source volume, except
       for the files with no extension (such as README), or files
       with extension .doc, .wpd, .ppt. Then exclude any files in
       \temp directory, except for files under \temp\system
       directory, or named \temp\download.lst, or matches wildcard
       pattern \temp\import*.exe (such as \temp\import1.exe).

  -->

    <exclude>

      <path>\</path>

    </exclude>

    <include>

      <extension>.</extension>       <!-- Path with no extension -->

      <extension>.doc</extension>
```

```
        <extension>.wpd</extension>

        <extension>.ppt</extension>

      </include>

      <exclude>

        <path>\temp</path>

      </exclude>

      <include>

        <path>\temp\system</path>

        <path>\temp\download.lst</path>

        <wildcard>\temp\import*.exe</wildcard>

      </include>

    </filter>

  </source>

</job> <!-- End Job element for server1_volume1 -->

<job> <!-- Begin another Job element for server2_volume2. -->
```

```
<name>server2_volume2</name>

<stopped/> <!-- The job is in a Stopped state. -->

<frequency>

  <interval unit="minutes">15</interval>

</frequency>

<source>

  <server>server2</server>

  <!-- &, <, and > are special characters in XML, and must be
       enclosed by CDATA. In this example, volume2& must be
       enclosed in the CDATA command.
  -->

  <volume><![CDATA[volume2&]]></volume>

</source>

<deletePolicy> <!-- How many and how long to keep versions in the
                    archive -->

  <interval unit="days">2</interval>

  <maxKeepVersions>2</maxKeepVersions>

  <maxKeepTime unit="days">90</maxKeepTime>
```

```
    </deletePolicy>


  </job>  <!-- End another Job element -->


  <!-- Add a Job element for every volume that has data you want to
       version. -->


</arkConfig> <!-- End Root element -->


<!-- End of arkConfig.xml -->
```

Troubleshooting

C

This section discusses potential issues and workarounds for Novell® Archive and Version Services.

C.1 Archive Server

This section discusses known issues and workarounds for debugging the Archive server.

Versions are not available when user maps the Primary volume using CIFS

Possible Cause: The Primary volume is mapped through CIFS using IP address.

Action: To resolve this issue, map the Primary volume through CIFS using the hostname.

C.2 Upgrade Issues

This section discusses known issues and workarounds for upgrading.

Missing Parameter in the ArkSQL.cnf File after Upgrading from NetWare 6.5 SP1

Problem: The Transaction Isolation parameter in the `sys:\arkManager\arkSQL.cnf` file is missing or is set to the wrong value after you upgrade from NetWare 6.5 SP1 to any later versions of NetWare. Data corruption is possible.

Possible Cause: The Transaction Isolation parameter in the `sys:\arkManager\arkSQL.cnf` file is not set to READ_COMMITTED or is missing from the file.

Action: If you upgrade directly from NetWare 6.5 SP1 to any later version of NetWare, you must modify the `sys:\arkManager\arkSQL.cnf` file. To prevent possible data corruption, set the Transaction Isolation parameter to Read_Committed before restarting ArkManager after the upgrade.

1. Stop ArkManager 2.0.
2. Upgrade the server to OES NetWare.
3. Verify that the server is operating as expected, but do not start ArkManager.
4. Compare and modify values of non-system-specific information in `sys:\arkManager\arkSQL.cnf` with the ones defined in `sys:\arkManager\arkSQL_sample.cnf`.

   Most importantly, make sure to add the following variable setting to the end of the file:

```
set-variable = transaction_isolation=READ_COMMITTED
```

5 Start ArkManager 2.0.

If you are configuring a new archive server after the upgrade, this setting is part of the updated sample `arkSQL.cnf` file, `sys:\arkManager\arkSQL_sample.cnf`. Use the updated sample file as a guide when setting up additional archive servers.

IMPORTANT: This defect exists for upgrades from NetWare SP1 to NetWare SP2 or later servers or OES NetWare servers. It was resolved in NetWare SP2 and is not a problem for upgrades from NetWare SP2.

C.3 MySQL Issues

This section discusses known issues and workarounds for the MySQL database.

Possible File Lock Conflicts in the Archive Database

Problem: Users experience errors when deleting or restoring file versions.

Possible Cause: The errors might be caused by file lock conflicts in the archive database.

Action: To disable database locks:

1 Stop ArkManager by entering the following at the server console:

```
arkstop
```

2 Stop the MySQL server by entering the following at the server console:

```
mysqladmin -p shutdown --port=value
```

Replace *value* with the port number where the ArkManager instance of MySQL server is running, such as 3308.

3 Disable database locks by opening the `sys:\arkManager\arkSQL.cnf` file in a text editor and adding the following line to the end of the file:

```
set-variable = innodb_table_locks=OFF
```

4 Restart the MySQL server by entering the following at the server console (all on the same line, of course):

```
mysqld_safe --defaults-file= sys:\arkmanager\arksql.cnf
```

5 Restart ArkManager by entering the following at the server console:

```
arkstart
```

Documentation Updates

D

This section contains information about documentation content changes made to the *NW 6.5 SP8: Novell Archive and Version Services 2.0 Administration Guide* since the initial release of NetWare® 6.5 SP3. If you are an existing user, review the change entries to readily identify modified content. If you are a new user, simply read the guide in its current state.

This document was updated on the following dates:

D.1 November 9, 2009

This guide has been modified for publication on the NetWare 6.5 SP8 Documentation Web site.

D.2 October 2009

Updates were made to the following section. The changes are explained below.

D.2.1 Installing and Configuring an Archive Server Cluster

Location	Change
Installing and Configuring an Archive Server Cluster	Moved the section and created a new chapter.

D.3 January 1, 2007 (NetWare 6.5 SP6)

Updates were made to the following sections. The changes are explained below.

D.3.1 Preface

The following change was made to this section:

Location	Change
Additional Documentation	For information about in the command files for ArkManager, see "Archive Definitions" (http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/bsin3bl.html) in the *NDK: Virtual File Services for NetWare* (http://developer.novell.com/ndk/vfs.htm).

D.3.2 Coexistence and Migration Issues for Archive and Version Services

The following change was made to this section:

Location	Change
Interoperability of the Archive Versioning Plug-In	This section is new.

D.4 November 1, 2005

The entire guide was reformatted to comply with revised Novell documentation standards. The content is unchanged.

D.5 August 19, 2005 (NetWare 6.5 SP4)

Updates were made to the following sections. The changes are explained below.

- Section D.5.1, "Coexistence and Migration Issues for Archive and Version Services," on page 158
- Section D.5.2, "Configuring Jobs in iManager," on page 158
- Section D.5.3, "Managing Jobs," on page 158
- Section D.5.4, "Security Considerations for Archive and Version Services," on page 159

D.5.1 Coexistence and Migration Issues for Archive and Version Services

This section is new. It contains information that previously appeared in Prerequisites and Guidelines.

D.5.2 Configuring Jobs in iManager

This section is new and is based on the redesigned Archive Versioning plug-in for iManager 2.5.

D.5.3 Managing Jobs

This section is new and is based on the redesigned Archive Versioning plug-in for iManager 2.5.

D.5.4 Security Considerations for Archive and Version Services

This section is new. It contains information that previously appeared in Prerequisites and Guidelines.